Received: Added at production | Revised: Added at production | Accepted: Added at production

DOI: xxx/xxxx

ARTICLE TYPE

Untamed: Unconstrained Tensor Decomposition and Graph Node Embedding for Cortical Parcellation

Yijun Liu[1] | **Jian Li**[2,3] | **Jessica L. Wisnowski**[4,5] | **Richard M. Leahy**[1]

[1]Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, CA, United States

[2]Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital and Harvard Medical School, Charlestown, United States

[3]Center for Neurotechnology and Neurorecovery, Department of Neurology, Massachusetts General Hospital and Harvard Medical School, MA, United States

[4]Radiology and Pediatrics, Division of Neonatology, Children's Hospital Los Angeles, CA, United States

[5]Keck School of Medicine, University of Southern California, CA, United States

Correspondence

Corresponding author: Richard M. Leahy
Email: leahy@sipi.usc.edu

Abstract

Cortical parcellation is fundamental to neuroscience, enabling the division of cerebral cortex into distinct, non-overlapping regions to support interpretation and comparison of complex neuroimaging data. Although extensive literature has investigated cortical parcellation and its connection to functional brain networks, the optimal spatial features for deriving parcellations from resting-state fMRI (rsfMRI) remain unclear. Traditional methods such as Independent Component Analysis (ICA) have been widely used to identify large-scale functional networks, while other approaches define disjoint cortical parcellations. However, bridging these perspectives through effective feature extraction remains an open challenge. To address this, we introduce *Untamed*, a novel framework that integrates unconstrained tensor decomposition using NASCAR to identify functional networks, with state-of-the-art graph node embedding to generate cortical parcellations. Our method produces near-homogeneous, spatially coherent regions aligned with large-scale functional networks, while avoiding strong assumptions like statistical independence required in ICA. Across multiple datasets, Untamed consistently demonstrates improved or comparable performance in functional connectivity homogeneity and task contrast alignment compared to existing atlases. The pipeline is fully automated, allowing for rapid adaptation to new datasets and the generation of custom parcellations. The atlases derived from the Genomics Superstruct Project (GSP) dataset, along with the code for generating customizable parcel numbers, are publicly available at https://untamed-atlas.github.io.

KEYWORDS

Cortical parcellation, Resting-state fMRI, Temporal synchronization, Tensor decomposition, Graph representation learning

1 | INTRODUCTION

Elucidating the macrostructure of the human brain remains a cornerstone in neuroscience. From the early work of Brodmann though Talairach through the most recent efforts (Brodmann, 1909; Tzourio-Mazoyer et al., 2002; Schaefer et al., 2018; Eickhoff, Yeo, & Genon, 2018; Han et al., 2020), a key aim is to identify the parcels – the relatively homogeneous mesoscale computational units that together give rise to the large-scale neural networks of the human brain.

The most common approach to elucidating the intrinsic macrostructure involves resting-state functional magnetic resonance imaging (rsfMRI). By recording blood-oxygen-level-dependent (BOLD) signals without task stimuli, rsfMRI proved instrumental in deciphering the human brain's intrinsic functional brain networks (Allen et al., 2014; Beckmann, DeLuca, Devlin, & Smith, 2005; Biswal, Yetkin, Haughton, & Hyde, 1995). From these, investigators have derived both group-level (Yeo et al., 2011; Craddock, James, Holtzheimer, Hu, & Mayberg, 2012; Shen, Tokoglu, Papademetris, & Constable, 2013; Gordon et al., 2016; Schaefer et al., 2018; Yan et al., 2023; Peng et al., 2023), and individual (Chong et al., 2017; Kong et al., 2019) parcellations of the human brain. Until recently, resting-state networks were often derived under assumptions of statistical independence or orthogonality, as imposed by independent component analysis (ICA) and principal component analysis (PCA). However, large-scale brain systems are known to interact extensively, both anatomically and functionally. For instance, the default mode network and task-positive networks exhibit systematic anti-correlations in resting-state fMRI (Fox et al., 2005), indicating that intrinsic networks are functionally coupled rather than statistically independent. Consequently, enforcing independence or orthogonality can artificially separate signals that are, in reality, partially correlated. Such constraints may therefore fail to capture the true organization of intrinsic functional brain networks (Harris & Mrsic-Flogel, 2013; Rockland, 2019; Hutchison et al., 2013; Karahanoğlu & Van De Ville, 2015; Seeley et al., 2007; Leech, Braga, & Sharp, 2012).

We have developed a new way for identifying overlapping and non-orthogonal functional networks using a Nadam-Accelerated SCAlabe and Robust (NASCAR) tensor decomposition method (Li, Wisnowski, Joshi, & Leahy, 2021; Li, Liu, Wisnowski, & Leahy, 2023). The aim of this study is to determine whether cortical parcellations, derived from NASCAR networks, would be more robust when compared to prior parcellations. Briefly, NASCAR computes a three-way tensor decomposition of temporally aligned rsfMRI data across subjects (Joshi, Chong, Li, Choi, & Leahy, 2018; Akrami, Joshi, Li, & Leahy, 2019) without imposing implausible constraints, such as statistical independence or orthogonality (Smith, Hyvärinen, Varoquaux, Miller, & Beckmann, 2014). It facilitates the identification of functional networks that may more authentically represent brain activity patterns by avoiding the need to impose spatial and/or temporal independence between networks. The networks identified by NASCAR have been shown to be, to a large degree, consistent with those found using ICA but with subtle differences that are consistent with other findings in the literature, such as subsystems of the default-mode network found using seed-based clustering and Bayesian methods (Andrews-Hanna, Reidler, Sepulcre, Poulin, & Buckner, 2010; Andrews-Hanna, Smallwood, & Spreng, 2014; Buckner & DiNicola, 2019; Harrison et al., 2020). We also showed that the networks found using NASCAR could be used to better classify subjects with neurological disorders (Li et al., 2023).

Based on the brain networks identified by NASCAR, we construct a graph and apply a graph node embedding method, NetMF (Qiu et al., 2018), to generate cortical parcellations. This integration of tensor decomposition and graph node embedding enables the production of parcellations that are functionally homogeneous, computationally efficient, and adaptable to a wide range of datasets. Unlike parcellation methods that rely on manual input (e.g., (Glasser et al., 2016; Joshi et al., 2022)), our pipeline is fully automated, enabling researchers to seamlessly apply the framework to new datasets with any desired number of parcels. This automation facilitates broader exploration of neuroscience questions, such as examining the influence of functional network spatial maps on the resulting parcels and determining the optimal number of parcels, as demonstrated in our experiments.

Our approach, which we term *Untamed* (**Un**constrained **t**ensor decomposition-based **a**ctivation **m**aps and emb**ed**ding), delivers performance that exceeds or is comparable to the state-of-the-art parcellations (e.g., (Schaefer et al., 2018)) across two critical metrics: resting-state functional connectivity homogeneity and task contrast alignment. By addressing methodological challenges in determining spatial features and linking cortical parcellation with functional networks, Untamed provides a robust framework for advancing research in neuroscientific studies. The parcellation label maps derived from the Genomics Superstruct Project (Holmes et al., 2015) dataset and the Human Connectome Project (Van Essen et al., 2012) and the corresponding code/scripts to generate those maps are available publicly at `https://untamed-atlas.github.io`.

2 | MATERIAL AND METHODS

2.1 | Overview

We derived cortical parcellations using rsfMRI data from 1428 subjects from the Genomics Superstruct Project (GSP) dataset (Holmes et al., 2015). The overall pipeline is depicted in Fig. 1. The procedure included a temporal alignment using BrainSync (Akrami et al., 2019; Joshi et al., 2018), a 3-way tensor decomposition using NASCAR (Li et al., 2021, 2023), graph construction from spatial maps followed by graph node embedding using NetMF (Qiu et al., 2018), and finally *k*-means clustering to obtain the parcellations. We evaluated the parcellation performance on the Yale rsfMRI (Lee et al., 2022), Human Connectome Project (HCP) (Van Essen et al., 2012; Glasser et al., 2013), and the multi-domain task battery (MDTB) (King, Hernandez-Castillo, Poldrack, Ivry, & Diedrichsen, 2019; Zhi, King, Hernandez-Castillo, & Diedrichsen, 2022) datasets, with varied acquisition protocols and data modalities.

2.2 | Datasets

2.2.1 | The Genomics Superstruct Project (GSP) dataset

We utilized the Genomics Superstruct Project (GSP) (Holmes et al., 2015) dataset for atlas generation. The dataset contains 3T resting-state fMRI (rsfMRI) scans from 1570 subjects (665 males, 905 females, age between 22 and 35). Each subject underwent either one or two rsfMRI scans. Of the 1570 participants, 1139 completed two scans, while 431 completed only a single scan. The rsfMRI data were acquired with *TR* = 3s, and a 3 mm isotropic spatial resolution. Each session lasted 6 min and 12 s, yielding 124 time points. The structural data associated with each subject consisted of a single 1.2 mm isotropic scan.

Data preprocesssing was conducted using the publicly available pipeline from the CBIG repository (`https://github.com/ThomasYeoLab/Standalone_CBIG_fMRI_Preproc2016`), configured to closely align with methods described in (Schaefer et al., 2018; Li et al., 2019; Yan et al., 2023). Each subject's T1-weighted image was registered to the standard MNI 2mm space. Key preprocessing steps of fMRI data included slice time correction, motion correction, censoring of outlier volumes, signal regression of white matter, ventricular signals, whole brain signal, and bandpass filtering (0.009 Hz – 0.08 Hz). Alignment between fMRI and structural images was performed using boundary-based registration (Greve & Fischl, 2009). While we acknowledge that alternative surface-based alignment approaches may yield improved alignment performance (Coalson, Van Essen, & Glasser, 2018), we adopted the same volume-based alignment strategy used in (Schaefer et al., 2018; Yan et al., 2023) to ensure that comparisons with these widely used atlases were conducted under matched conditions.

The preprocessed fMRI were projected to the Freesurfer's fsaverage6 space and smoothed using a Gaussian kernel with a 6 mm full-width half-maximum (FWHM). Subjects with at least one rsfMRI session that passed quality control checks in the preprocessing pipeline were included in subsequent experiments. This resulted in a final sample of 1428 subjects (603 males, 825 females), of which 1034 had two scans and 394 had one scan.

2.2.2 | The Human Connectome Project (HCP) dataset

We utilized the minimally preprocessed 3T rsfMRI data from 1000 subjects (466 males, 534 females, age between 22 and 35) from the Human Connec tome Project (HCP) (Glasser et al., 2013; Van Essen et al., 2012) for the assessment of parcellations. The rsfMRI data were acquired with *TR* = 0.72 s, *TE* = 33.1 ms, and a 2 mm isotropic resolution and co-registered onto a common atlas in MNI space (Glasser et al., 2013). We used the scans acquired in the LR phase encoding direction. Each session ran 15 minutes with 1200 time points. The data were resampled onto the cortical surface extracted from each subject's T1-weighted MRI and co-registered to a common surface (Glasser et al., 2013). No additional spatial smoothing was applied other than the 2 mm FWHM isotropic Gaussian smoothing in the minimal preprocessing pipeline (Glasser et al., 2013).

In addition to rsfMRI data, we utilized subject-wise task activation z-score maps in fs_LR 32K surface space from the HCP dataset (Barch et al., 2013) for evaluation. These maps covered seven task domains: working memory, gambling, motor, language, social, relational, and emotion. Task contrast maps were derived from 3T task fMRI data acquired with the same parameters as the rsfMRI data, except for the duration which varied based on the particular task. We included all available contrasts and incorporated data from 1,006 subjects with complete task datasets. Furthermore, group-level task activation z-score maps for the same tasks were also employed for the analysis of optimal parcel numbers.

2.2.3 | Yale resting-state fMRI dataset

We used the Yale rsfMRI dataset (Lee et al., 2022) as an independent dataset for assessing RSFC homogeneity (Schaefer et al., 2018). The dataset comprises 27 subjects (11 males, 16 females, aged between 22 and 31). Each subject underwent two rsfMRI scans, each with *TR* = 1s, *TE* = 30ms, and a 2 mm isotropic resolution. The total duration of each session was 6 minutes 50 seconds, encompassing 410 frames. We preprocessed the data using fMRIPrep (Esteban et al., 2019), which sampled the fMRI data onto the fs_LR 32K surfaces compatible with the HCP dataset. Four subjects' data were not successfully preprocessed by the fMRIPrep pipeline due to corrupted files and were consequently excluded from the evaluation. Following the recommendations from the curators of the dataset, we discarded the first 10 seconds of each scan and temporally concatenated the rsfMRI data of the two sessions for each subject.

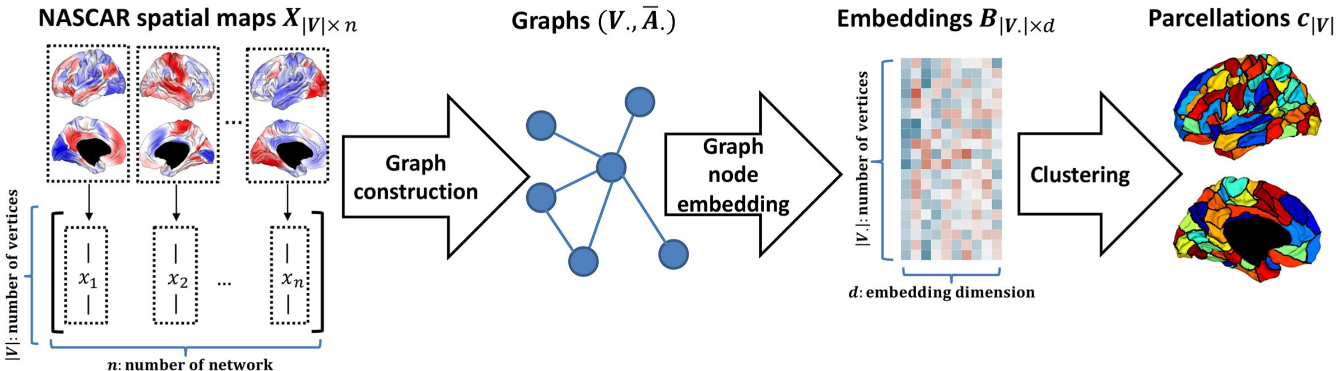

FIGURE 1 Overview of the Untamed parcellation framework. NASCAR tensor decomposition yields spatial network maps $\boldsymbol{X} \in \mathbb{R}^{|V|\times n}$, where $|\boldsymbol{V}|$ is the number of cortical vertices and n is the number of functional networks. For each hemisphere, a graph $(\boldsymbol{V}_., \bar{\boldsymbol{A}}_.)$ is constructed based on the pairwise correlations of network participation across vertices. Graph node embeddings $\boldsymbol{B} \in \mathbb{R}^{|\boldsymbol{V}_.|\times d}$ are derived using NetMF, producing d-dimensional representations for each vertex. Finally, clustering in the embedding space yields parcellations $\boldsymbol{c}_{|V|}$, i.e., assignments of vertices into parcels at the chosen resolution.

2.2.4 | Multi-domain task battery (MDTB) dataset

We employed the multi-domain task battery (MDTB) dataset (King et al., 2019; Zhi et al., 2022) for evaluating the alignment between parcels in the atlases and areas of activation in the task activation maps. The MDTB dataset contains task fMRI for 24 healthy subjects (8 males, 16 females, mean age 23.8 years old) conducting 26 tasks, including motor, language, and social domains. The fMRI data were acquired on a 3T Siemens Prisma scanner with $TR = 1$s, 3 mm slice thickness, and $2.5 \times 2.5\ mm^2$ in-plane resolution. The contrast maps were derived using general linear modeling (GLM) based on the task designs and are available from the database, already resampled to the fs_LR 32K space.

2.3 | Tensor-based identification of Brain Networks using NASCAR and BrainSync

Following the procedure in (Schaefer et al., 2018), when two resting-state fMRI sessions were available for a subject, the corresponding time series were concatenated temporally. When only a single session was available, that session was duplicated and concatenated to match the temporal length of subjects with two sessions, thereby avoiding variability in data length across subjects during downstream processing. The resulting multi-subject rsfMRI data were then organized as a third-order tensor (space × time × subject). In order to identify a low-rank model via tensor decomposition, the data needed to be both temporally and spatially aligned. As described in the preprocessing pipeline for the GSP dataset, inter-subject spatial alignment was achieved using standard image registration methods, mapping each subject's data to the standard fsaverage6 surface space. However, while spatial alignment ensured that brain regions corresponded across subjects, the rsfMRI time series remained independent for each subject. To address this, we assume that subjects share a similar connectivity structure as reflected in the pairwise correlation patterns between brain regions. We note that this shared-structure assumption is common to most group-level functional parcellation frameworks, including clustering-, decomposition-, and gradient-based approaches (Yeo et al., 2011; Schaefer et al., 2018). In the present work, this assumption is used to enable group-level temporal synchronization and network estimation, rather than being imposed as a hard mathematical constraint on the spatial components. Unlike methods that enforce orthogonality or statistical independence, the NASCAR-based representation allows correlated and overlapping network structure to emerge from the data. We can then apply an orthogonal transformation to temporally align or synchronize the data across subjects. This temporal alignment is achieved using the BrainSync transform (Joshi et al., 2018). BrainSync computes an orthogonal transformation between fMRI recordings from a pair of subjects such that the sum of squared errors between their aligned time series is minimized. As a result, the time series in homologous brain locations will be highly correlated after the BrainSync transform is applied (and perfectly so if the two datasets have identical correlations). A multi-subject extension, as described in (Akrami et al., 2019), minimizes the squared error from all subjects to an automatically generated group template. This method was used in this study to align all subjects for generating the cortical parcellation. BrainSync applies an orthogonal and invertible transformation to temporally align subjects, which preserves second-order statistical structure up to rotation while reducing latency-related variability, without enforcing artificial similarity across subjects (Akrami et al., 2019; Li et al., 2023). After temporal synchronization, we formed a third-order tensor $\boldsymbol{\mathcal{X}} \in \mathbb{R}^{|\boldsymbol{V}| \times T \times S}$ from the spatially aligned and temporally synchronized rsfMRI data, where $\boldsymbol{V}$ is the set of cortical vertices with a cardinality $|\boldsymbol{V}| \approx 75K$, $T = 240$ is the number of time points, $S = 1428$ is the number of subjects used to generate the atlas. We then applied NASCAR to decompose the group rsfMRI data into a set of brain networks common to all subjects using a Canonical Polyadic model (Kolda & Bader, 2009; Cichocki et al., 2015; Li et al., 2021, 2023)

$$\boldsymbol{\mathcal{X}} \approx \sum_{i=1}^{R} \lambda_i \, \boldsymbol{a}_i \circ \boldsymbol{b}_i \circ \boldsymbol{c}_i, \tag{1}$$

where each outer product $\lambda_i \, \boldsymbol{a}_i \circ \boldsymbol{b}_i \circ \boldsymbol{c}_i$ represents a brain network: $\boldsymbol{a}_i \in \mathbb{R}^{|\boldsymbol{V}|}$ are the spatial network maps, $\boldsymbol{b}_i \in \mathbb{R}^{T}$ their (synchronized) temporal dynamics, and $\boldsymbol{c}_i \in \mathbb{R}^{S}$ the subject participation level for the i^{th} network; λ_i represents the network magnitude, indicating the relative activity level with respect to other networks. In contrast to the commonly used blind source separation techniques such as PCA or ICA, NASCAR imposes neither orthogonality (as with PCA) (Smith et al., 2014) nor statistical independence (as with ICA) (Beckmann et al., 2005). The shared spatial network maps $\{\boldsymbol{a}_i\}$ can be overlapped and correlated (Li et al., 2023). Based on our previous work that examined stability and reproducibility across data sets, we chose a rank R=50 (Li et al., 2023). We performed NASCAR decomposition on the whole GSP dataset and obtained a set of 50 spatial network maps, which formed a spatial feature matrix $\boldsymbol{X} \in \mathbb{R}^{|\boldsymbol{V}| \times \mathbf{50}}$ ($|\boldsymbol{V}| \approx 75K$). This matrix provided a 50-dimensional feature vector for each vertex. Specifically, $\boldsymbol{X}_{ij}$ represented the participation level at the i^{th} vertex in the j^{th} network map.

2.4 | Graph construction from NASCAR spatial maps

We computed the Pearson correlation $\boldsymbol{A} = \operatorname{corr}(\boldsymbol{X}) \in \mathbb{R}^{|\boldsymbol{V}| \times |\boldsymbol{V}|}$ from the feature matrix $\boldsymbol{X}$ as a measure of similarity between feature vectors of each pair of vertices. We then computed an adjacency matrix, following (Ng, Jordan, & Weiss, 2001), using the Gaussian kernel $\widetilde{\boldsymbol{A}}_{i,j} = \exp\left(\frac{\boldsymbol{A}_{i,j}}{2\sigma^2}\right)$, $\sigma = 0.5$. Our goal was to produce contiguous parcels, and as such, we were not interested in long-range correlations such as inter-hemispheric connections. Therefore, following the practice in (Schaefer et al., 2018), we created two separate graphs for the left and right hemispheres. Specifically, we defined $\overline{\boldsymbol{A}}_{i,j} = \widetilde{\boldsymbol{A}}_{i,j}$ for all (i,j) for which i is within nb hops of j and zero otherwise. In a related approach, (Craddock et al., 2012) used a 1-hop neighborhood, retaining nearest neighbors only. However, we found that a larger nb could produce higher RSFC homogeneity (see Evaluation section). Consequently, rather than fixing this parameter, we treated it as a hyperparameter within the algorithm, optimized using the GSP dataset.

Using the above approach, we obtained two sparsely connected (nb-hop) graphs for each hemisphere: $\boldsymbol{G}_L = (\boldsymbol{V}_L, \overline{\boldsymbol{A}}_L)$ and $\boldsymbol{G}_R = (\boldsymbol{V}_R, \overline{\boldsymbol{A}}_R)$, where $\boldsymbol{V}_L, \boldsymbol{V}_R$ were the set of cortical vertices and $\boldsymbol{A}_L, \boldsymbol{A}_R$ were the nb-hop filtered adjacency matrices of the graphs $\boldsymbol{G}_L$, $\boldsymbol{G}_R$, for the left and the right hemispheres, respectively.

2.5 | Graph node embedding and clustering

We adopted the method outlined in NetMF (Qiu et al., 2018) to embed vertices in a lower dimensional space with more representative features that can be used for clustering. This approach (Qiu et al., 2018) established an equivalence between the widely-used DeepWalk algorithm (Perozzi, Al-Rfou, & Skiena, 2014) and matrix factorization techniques. Briefly, DeepWalk traverses vertices in the graph using a random walk. At each vertex during traversal, DeepWalk considers neighboring vertices as co-occurrent "positive" pairs as well as random node pairs outside the neighborhood as "negative samples." It then applies a skip-gram model to train on the collected samples to derive the final embeddings. NetMF approximates DeepWalk by factorizing the following matrix (Qiu et al., 2018):

$$\boldsymbol{M} = \log\left(\max\left\{\frac{\sum_{i=1}^{|\boldsymbol{V}|}\sum_{j=1}^{|\boldsymbol{V}|} \overline{\boldsymbol{A}}_{i,j}}{bT} \sum_{r=1}^{T} \left(\boldsymbol{D}^{-1}\overline{\boldsymbol{A}}\right)^{r} \boldsymbol{D}^{-1},\ 1\ \right\}\right), \tag{2}$$

where $\boldsymbol{D}$ is the degree matrix of the graph; $\log(\cdot)$ and $\max\{\cdot, 1\}$ are element-wise logarithmic and maximum operations, respectively; T is the window length over which nodes considered as co-occurrent positive pairs; and b denotes the expected number of unrelated node pairs each positive pair is contrasted against when learning its embedding. We found that in practice

TABLE 1 Embedding and neighborhood hyperparameters by parcel count range.

Parcel range	Embedding dim.	Neighborhood size (*nb*)	Window length (*T*)
(0, 200]	128	55	7
(200, 300]	128	40	15
(300, 400]	128	35	15

the method was not sensitive to the choice of these two hyperparameters, which had little impact on the averaged training subjects' RSFC homogeneity (see evaluation metrics section). We set $b = 1$, the default value in the original NetMF paper. We computed the $\boldsymbol{M}$ matrices for left hemisphere from $\overline{\boldsymbol{A}}_L$ and for right hemisphere from $\overline{\boldsymbol{A}}_R$ separately. We used singular value decomposition (SVD) to obtain the final embeddings. Specifically, let $\boldsymbol{M} = \boldsymbol{U}_d \boldsymbol{\Sigma}_d \boldsymbol{V}_d^{\mathrm{T}}$, where $\boldsymbol{\Sigma}_d$ and $\boldsymbol{U}_d$ being the d largest singular values and their associated left singular vectors. The final embeddings were computed as $\boldsymbol{B} = \boldsymbol{U}_d \sqrt{\boldsymbol{\Sigma}_d}$ for the left $(\boldsymbol{B}_L)$ and right hemisphere $(\boldsymbol{B}_R)$ separately, with d being the embedding dimension.

To obtain the final cortical parcellations, we applied k-means clustering separately to the embedding vectors corresponding to the left $\boldsymbol{B}_L$ and right hemispheres $\boldsymbol{B}_R$ hemispheres. The number of clusters *k* was set to match the number of desired parcels per hemisphere. For each k, we ran k-means clustering with 500 different random initializations with a maximum of 20000 iterations and selected the result that yielded the smallest cost across trials as the final output. This clustering step produces a discrete parcellation in which each vertex is assigned to exactly one parcel. While the underlying NASCAR-derived network representations may exhibit correlated or partially overlapping patterns of network expression, overlap is not permitted at the level of parcel assignment, ensuring that the final parcellation consists of spatially disjoint regions.

The graph construction and embedding procedure involved fine tuning the following parameters: spatial neighborhood constraint nb, the window length T, and the embedding dimension d. Specifically, nb was varied across the set of $\{1, 5, 10, 15, 20, \ldots, 60\}$, T across $\{1, 7, 11, 15\}$, and d across $\{128, 256\}$. The fine tuning process aimed to maximize the weighted average RSFC homogeneity. Hyperparameter combinations were selected based on the mean weighted average RSFC homogeneity across GSP subjects, which served as a summary statistic to assess performance on the training set.

The parameter combinations used in the Untamed framework were selected based on the parameter sweeping results shown in Table S2 and the final parameter combinations as shown in Table 1.

2.6 | Evaluation metrics

2.6.1 | Resting-state functional connectivity (RSFC) homogeneity

Each subject's RSFC was computed as the Pearson correlation of rsfMRI time series between all pairs of cortical vertices. A Fisher z-transform was subsequently performed to obtain z scores from the correlation values. Similar to the evaluation procedure in (Joshi et al., 2022; Schaefer et al., 2018), a parcel-wise homogeneity score ρ_i was computed as the averaged RSFC values within the i^{th} parcel. To obtain a global homogeneity measure for a single subject, a weighted average ρ of each parcel's homogeneity scores was then computed by accounting for different cluster sizes:

$$\rho = \sum_{i=1}^{N} \rho_i \frac{|\boldsymbol{V}^i|}{|\boldsymbol{V}|}, \tag{3}$$

where $|\boldsymbol{V}^i|$ is the number of vertices in parcel i, and *N* is the total number of parcels; $|\boldsymbol{V}|$ is the total number of cortical vertices. For conciseness, we refer to this RSFC weighted average homogeneity as "homogeneity" hereafter. The homogeneity was computed for each test subject separately. To quantitatively compare the homogeneity of Untamed to that of other atlases, we performed a paired-sample t-test. Effect size was reported using Cohen's *d* measure.

2.6.2 | Alignment with task contrasts

The degree to which the regions delineated in a particular parcellation reflect functional specialization was assessed by computing the variance of the task contrast within each parcel. The better the parcellation delineates regions of functional homogeneity, the lower the variance of task contrasts within each parcel. Similar to the procedure in (Schaefer et al., 2018; Joshi et al., 2022;

Yan et al., 2023), the variance metric was first computed for each parcel, and then a weighted average was computed, accounting for parcel size differences, as:

$$\sigma^2 = \sum_{i=1}^{N} \sigma_i^2 \frac{|V^i|}{|V|}, \qquad (4)$$

We refer to this weighted average task contrast variance as "task variance" hereafter. The variance was computed separately for each task contrast and for each subject. These variances were first averaged across contrasts within each task and then averaged across tasks, consistent with the approach in (Yan et al., 2023). Statistical significance was assessed using paired-sample t-test to compare the task variance between different parcellations.

2.7 | Comparison with existing parcellations

We compared our Untamed parcellations to an extensive set of 14 atlases listed in Table 2. Because comparison of the evaluation metrics defined above was only meaningful when comparing parcellations of the same parcellation resolution, in each case, we matched the number of parcels found using Untamed to the number in the left and right hemispheres for each baseline comparison.

The evaluations were conducted in the HCP fs_LR 32K surface space. Untamed, originally constructed in the fsaverage 6 surface space, was projected onto the HCP fs_LR 32K space. For other atlases that were originally constructed in spaces different from HCP fs_LR 32K (e.g., MNI152, fsaverage), we utilized versions of the atlases that had been resampled onto the HCP fs_LR 32K surfaces. These resampled versions were provided either by the original authors or by third parties, as detailed in Table 2. For the original two atlases proposed by (Yeo et al., 2011) that are spatially distributed across hemispheres with a relatively small number (7 and 17) of networks, we used the version provided in their GitHub repository (Table 2), where the distributed spatial networks were decomposed into local contiguous parcels (51 parcels for the 7 networks and 114 parcels for the 17 networks).

TABLE 2: List of atlases included in the comparison. "Functional": built from fMRI; "Hybrid": built from both fMRI and anatomical information.

Name	Type	# Clusters	Reference	Atlas Source
Yeo-51/114	Functional	51 (L: 26, R: 25) 114 (L: 57, R: 57)	(Yeo et al., 2011)	https://github.com/ThomasYeoLab/CBIG/tree/master/stable_projects/brain_parcellation
Power	Functional	130 (L: 65, R: 65)	(Arslan et al., 2018)	https://biomedia.doc.ic.ac.uk/brain-parcellation-survey/
USCBrain	Hybrid	130 (L: 65, R: 65)	(Joshi et al., 2022)	https://github.com/ajoshiusc/bfp/blob/main/supp_data/USCBrain_grayordinate_labels_clean.mat
Shen	Functional	200 (L: 102, R: 98)	(Shen et al., 2013)	https://biomedia.doc.ic.ac.uk/brain-parcellation-survey/
Schaefer-100/200/300/400	Functional	100 (L: 50, R: 50) 200 (L: 100, R: 100) 300 (L: 150, R: 150) 400 (L: 200, R: 200)	(Schaefer et al., 2018)	https://github.com/ThomasYeoLab/CBIG/tree/master/stable_projects/brain_parcellation
Yan-100/200/300/400	Functional	100 (L: 50, R: 50) 200 (L: 100, R: 100) 300 (L: 150, R: 150) 400 (L: 200, R: 200)	(Yan et al., 2023)	https://github.com/ThomasYeoLab/CBIG/tree/master/stable_projects/brain_parcellation
Glasser	Hybrid	360 (L: 180, R: 180)	(Glasser et al., 2016)	https://balsa.wustl.edu/study/show/RVVG

2.8 | Ablation study of input spatial maps: NASCAR versus ICA

In the Untamed framework, the graph was constructed using correlations between features representing the participation of each vertex in the NASCAR spatial network maps. However, the input spatial maps are not limited to NASCAR and can be derived from any network decomposition method. To investigate whether the non-orthogonal spatial maps generated by NASCAR lead to better performance compared to spatial maps from other methods, all steps in the pipeline were kept consistent, with the only variation being the input spatial maps.

For this study, ICA spatial maps provided by HCP were used as an alternative to NASCAR maps. Both the ICA and NASCAR maps consisted of 50 networks, with no additional selection or preprocessing applied. The pipeline's hyperparameters—including the number of spatial neighborhood constraints, the window length in the NetMF algorithm, and the embedding dimension—were optimized specifically for the ICA maps and explored within the same range as those used for NASCAR. RSFC homogeneity was evaluated on the HCP rsfMRI dataset consisting of 1000 test subjects using the same procedure described in Section 3.1. To statistically compare the homogeneity values derived from the NASCAR and ICA spatial maps, a paired-sample t-test was conducted. This analysis was designed to evaluate how the choice of input spatial maps influences the performance of the Untamed framework.

We note that the NASCAR spatial maps that Untamed atlases generated from were derived from the independent GSP dataset in the fsaverage 6 surface space and subsequently projected onto the HCP fs_LR 32K surface space. In contrast, the ICA maps from HCP inherently benefit from dataset-specific fine-tuning and remain within the same surface space, avoiding potential information loss associated with the projection process.

2.9 | Ablation study of graph embedding methods

We also compared Untamed with the spectral clustering described in (Ng et al., 2001), which applies a *k*-means clustering to the most significant eigenvectors of the normalized graph Laplacian. All steps in the learning procedure were identical except for the graph embedding, where we used NetMF. We applied both methods to the same graph constructed from NASCAR spatial maps as described in Section 2.4. Parcellations generated from NetMF and graph Laplacian were subject to the same neighborhood constraints, embedding dimension, and k-means clustering hyperparameters with a range of cluster numbers. We evaluated weighted average homogeneity on the 1000 HCP rsfMRI dataset test subjects. A paired-sample t-test test was used to compare the homogeneity values obtained from the Laplacian eigenvectors used in spectral clustering and the NetMF method that was utilized in Untamed.

2.10 | Ablation study of graph construction methods

In Untamed, we constructed the graph using the correlation between features representing the participation of each vertex in each of the NASCAR spatial network maps. Here, we conducted a comparison to the graph constructed from the widely used correlation map of RSFC, i.e. the correlation of the correlation between rsfMRI time series. Other steps were identical except for the construction of the graph. Parcellations generated from NASCAR spatial maps and RSFC correlation maps were subject to the same neighborhood constraints and embedding dimension. We evaluated weighted average homogeneity on the averaged RSFC from the 500 training subjects in the HCP dataset. A paired-sample t-test was employed to statistically compare homogeneity values obtained from NASCAR spatial maps and the correlation of RSFC.

2.11 | Is there a (natural) optimal number of parcels that can be identified from rsfMRI data

To address this question, we evaluated the performance of the Untamed atlas in identifying regions of relatively homogeneous functional activity, as measured by RSFC homogeneity (described in Section 3.1) and task contrast variance (described in Section 3.2), relative to random parcellations with an equal number of parcels on the HCP dataset. Random parcellations were employed as a null model to assess whether there was an optimal number of parcels for functional parcellation.

Specifically, we calculated the ratio of the metric scores obtained with the Untamed atlas to those from random parcellations for both RSFC homogeneity and task variance. Random parcellations were generated using a region-growing process implemented in the MNE-Python package (Gramfort, 2013). This process was repeated across a range of parcel numbers, from 1 to 500 parcels per hemisphere.

For each parcel number, 50 random parcellation realizations were generated with different initialization seeds. The mean weighted average metric scores across these 50 random parcellations were computed and used to calculate the performance ratio with respect to the Untamed atlas. These ratios were then plotted as a function of the number of parcels to investigate the relationship between parcel number and the performance of the Untamed atlas relative to the null model.

2.12 | RSFC-based parcellation characterization

To complement the weighted-average RSFC homogeneity metric that summarizes global homogeneity as a single value (Section 2.6.1), we further performed fine-grained analyses at the parcel and network levels.

2.12.1 | Parcel-wise RSFC analysis

This analysis examines how parcels differ locally in their RSFC homogeneity. For each subject, parcel-wise homogeneity was computed by averaging the rsfMRI time series within each parcel and calculating the mean Pearson correlation between the parcel-average time series and the time series of all vertices in that parcel. These values were then averaged across subjects. To compare atlases at matched resolution (for example Untamed-300 versus Schaefer-300), spatial correspondence between parcels was quantified using the Dice coefficient. Parcel-wise homogeneity profiles and Dice coefficients were plotted across parcel indices, and homogeneity values were visualized on the cortical surface for each atlas. Analogous analyses were conducted for the 100-parcel atlases, including comparisons with Random-100 generated as described in Section 2.11.

2.12.2 | Network-level RSFC analysis

This analysis examines how spatially contiguous parcels assemble into distributed functional networks. Using the Yale rsfMRI dataset, we first applied global signal regression and then averaged vertex time series within each parcel. Parcel-to-parcel RSFC matrices were computed using Pearson correlation and were Fisher-z transformed and averaged across subjects. For each atlas (Untamed-200 and Schaefer-200), networks were identified by applying spectral clustering (Ng et al., 2001) to the group-average RSFC matrix, with the number of clusters fixed at seven to match the canonical Yeo 7-network organization (Yeo et al., 2011). To relate the Untamed parcels to functional structure derived directly from NASCAR, parcel boundaries for Untamed-100 and Untamed-300 were overlaid on representative NASCAR spatial maps.

2.13 | NASCAR-derived network participation similarity

This analysis characterizes how similar vertices are in their patterns of participation across multiple functional networks. The NASCAR spatial maps obtained during graph construction (Section 2.4) provide, for each vertex, a vector describing its loading across the NASCAR-derived networks. We computed the spatial correlation between these loading vectors for all vertex pairs, yielding a similarity matrix that reflects how closely vertices match in their multi-network participation profiles. We kept the non-negative similarity values only and then computed, for each vertex, its “degree” as the sum of the non-negative similarities with all other vertices. We normalized the degrees to the range of $[0, 1]$, producing a cortical spatial map that quantifies how broadly its participation profile aligns with that of the rest of the cortex.

2.14 | Reproducibility analysis

We further quantify the similarity of parcel assignments using the Adjusted Rank Index (ARI) across groups within Untamed variants and between Untamed and other atlases (Fig. 11). ARI measures the agreement between two clustering solutions by considering all pairs of samples and counting the number of pairs assigned to the same or different clusters. For within-Untamed comparison, we randomly split the full set of GSP subjects into two equal subgroups (Group 1 and 2, each with 714 subjects) and generated separate atlases using the same pipeline described as in Sections 2.3–2.5. We then assessed reproducibility by computing ARI between the atlases from Group 1 and Group 2, and between Group 1 and that using the entire GSP dataset (the Untamed). For comparisons to other atlases, we similarly assessed the reproducibility using the ARI between Untamed, Schaefer, Yan, and Random atlases across parcel numbers.

3 | RESULTS

3.1 | Resting-state functional connectivity (RSFC) homogeneity

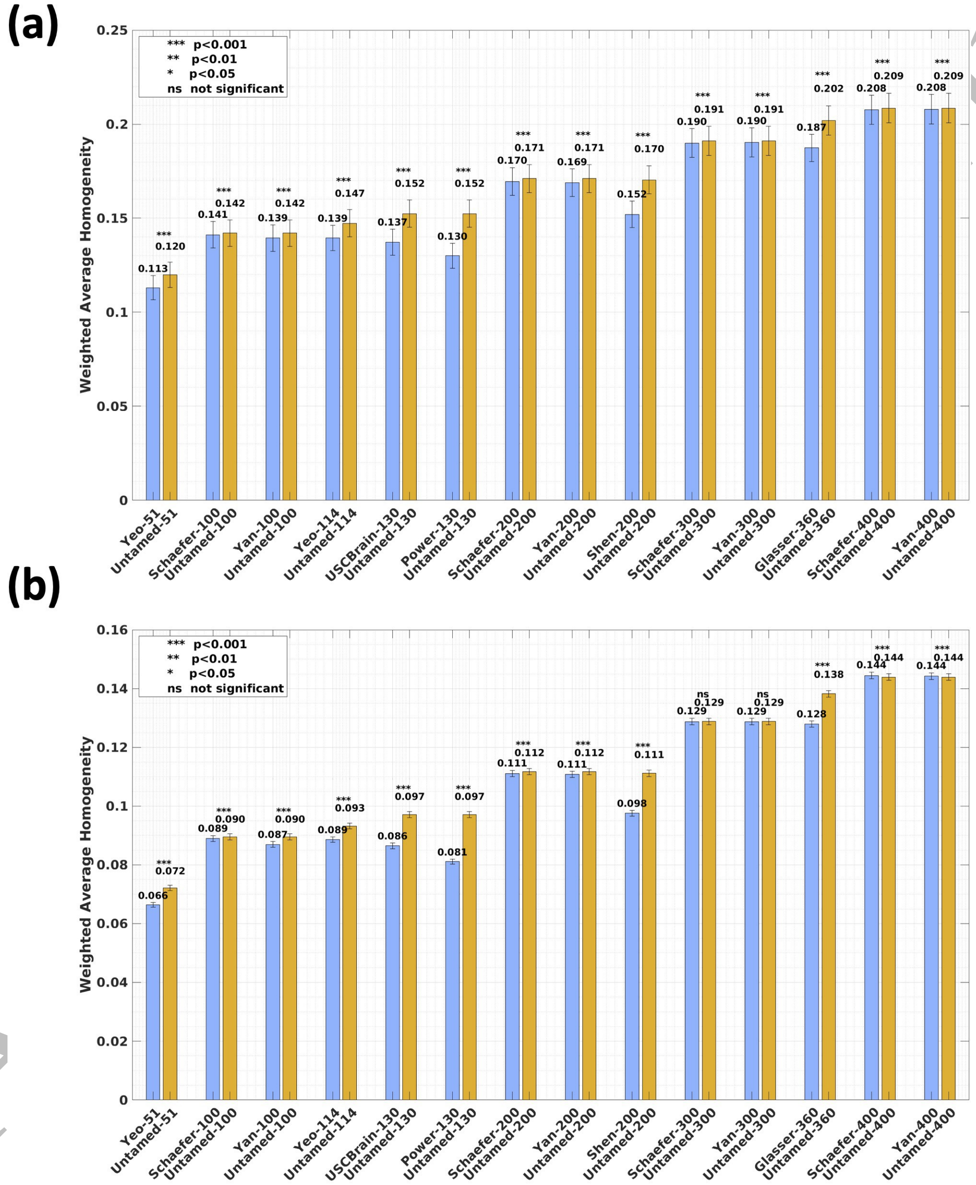

FIGURE 2 Weighted average RSFC homogeneity on the (a) Yale (b) HCP dataset. Each bar plot depicts the subject-wise RSFC weighted average homogeneity, averaged across test subjects, for each baseline and the Untamed atlas, with matched parcel numbers for the left and right hemispheres. The error bars represent the standard error across all subjects. Effect sizes are shown in Table S1.

Fig. 2 (a) (Yale dataset) and Fig. 2 (b) (HCP dataset) present the weighted average RSFC homogeneity as bar charts. Corresponding effect sizes, quantified by Cohen's d, are reported in Table S1. In Fig. 2 (a), results from the Yale dataset demonstrates that the Untamed method outperformed baseline methods across all tested parcel numbers with statistical significance ($p < .001$, *uncorrected*). Similarly, Fig. 2(b) shows that on the HCP dataset, the Untamed method consistently achieved higher weighted average homogeneity compared to baseline methods, with most differences being statistically significant ($p < .001$), except for the 300-parcel case, where no statistically significant differences were observed between Untamed-300 and Schaefer-300 or Yan-300.

An exception is noted in the 400-parcel case on the HCP dataset, where Schaefer-400 and Yan-400 achieve slightly higher homogeneity values than the Untamed method. However, the differences are negligible, with no variation up to the third decimal place and effect sizes of 0.0151 and 0.0085, respectively. Overall, these findings highlight the superior performance of the Untamed method, which generally produces more functionally homogeneous brain parcellations than other widely used approaches. This superiority was consistently observed across different datasets, parcel numbers, and parcellation schemes, emphasizing the robustness and effectiveness of the Untamed atlases. Interestingly, despite the large sample size of the HCP dataset, no statistical significance is detected between Untamed-300 and either Schaefer-300 or Yan-300. This suggests negligible differences between these methods in the 300-parcel case. This observation is further explored in the Discussion section.

3.2 | Alignment with task contrasts

The violin plots in Fig. 3 (a, b) depict the distribution of task contrast variance (as defined in Eq. 4) for the Untamed method and baseline atlases with matching parcel numbers. Fig. 3 (a) represents results from the MDTB task dataset, while Fig. 3 (b) illustrates results from the HCP task dataset. Lower task contrast variances correspond to a smaller overall contrast variance per parcel, indicating better alignment of parcellation with functional task responses.

Across the HCP dataset, the Untamed parcellations show statistically significant improvements over baseline atlases for nearly all matched-resolution comparisons. All significance assessments were based on paired-sample t-tests evaluated against the same Bonferroni-corrected threshold (α / number of parcellations), while raw p-values are reported for readability. The only comparison that does not reach the corrected threshold is Schaefer-300, where the task contrast variance is numerically lower for the baseline but the raw p-value ($p > .005$) is insufficient to pass the Bonferroni criterion. On the MDTB dataset, Schaefer-300 and Schaefer-400 yield slightly higher alignment scores than their Untamed counterparts, and these differences reach significance; however, the effect sizes are extremely small (differences indistinguishable at the third decimal place). For all remaining resolutions, Untamed either significantly outperforms the baselines or exhibits comparable performance (e.g., Schaefer-100/200 and Yan-300/400, where no significant differences are observed).

Although Glasser-360 explicitly used HCP task fMRI data (in combination with structural data) in constructing the atlas, the variances computed on both the HCP and MDTB task data were significantly higher than those of Untamed.

Fig. 4 (a, b) illustrates the overlay of Schaefer-100 and Untamed-100 parcels on the HCP group average "story" contrast map from the language task. Regions near Broca's area and other auditory regions demonstrate superior alignment of the Untamed parcels with contrast compared to Schaefer's parcellation. Conversely, a counterexample is presented in Figures 4 (c, d), where the "relational_match" contrast reveals better alignment of task-active regions with the Schaefer parcellation. The weighted average variance of language_story contrast is 4.3892 for Untamed-100 and 5.4311 for Schaefer-100. Meanwhile, the weighted average variance of relational_match contrast is 8.4673 for Untamed-100 and 8.0539 for Schaefer-100 (lower variance indicates better alignment). Overall, the Untamed method consistently achieves lower weighted average variance compared to baseline methods across most parcel numbers and parcellation schemes. This result highlights the ability of the Untamed method to produce functionally homogeneous brain regions with reduced variability in functional task responses.

3.3 | Ablation study of input spatial maps: NASCAR versus ICA

Fig. 5 shows the weighted average RSFC homogeneity of the 1000 subjects in the HCP dataset. The results indicates that despite the fact that using the ICA maps from HCP has the advantage of same-dataset hyperparameter fune-tining and remaining in the same surface space, using input spatial maps from NASCAR still consistently outperformed ICA across all four parcel numbers with statistical significance ($p < .001$).

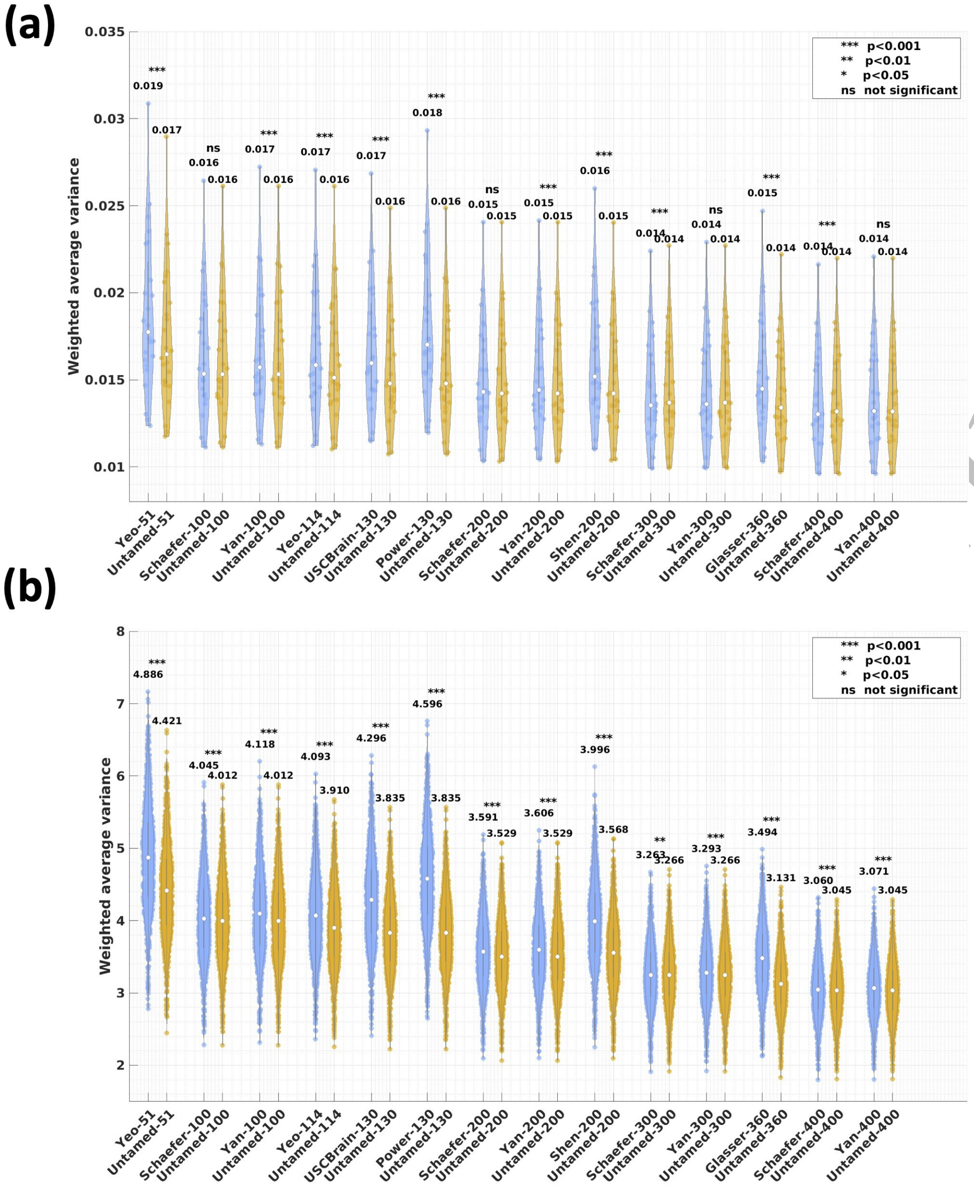

FIGURE 3 Weighted average task contrast variance evaluated on (a) MDTB (b) HCP test subjects. Each violin plot depicts the alignment with task contrast maps, computed for the baseline and the Untamed atlas, with matched parcel numbers for the left and right hemispheres.

3.4 | Ablation study of graph embedding methods

Supplementary Fig. S1 shows parcellations obtained using graph node embedding (NetMF) prior to clustering as described above with results obtained using spectral clustering directly from the eigenvectors of the graph Laplacian (GLC). Again, we show the bar plots of Untamed and GLC-based atlases in RSFC weighted average homogeneity per HCP test. Among the 4 different numbers of parcels tested, NetMF-based Untamed outperforms the GLC-based one in 3 cases with statistical significance ($p < .001$). These results support the use of NetMF embedding in place of the more standard GLC approach.

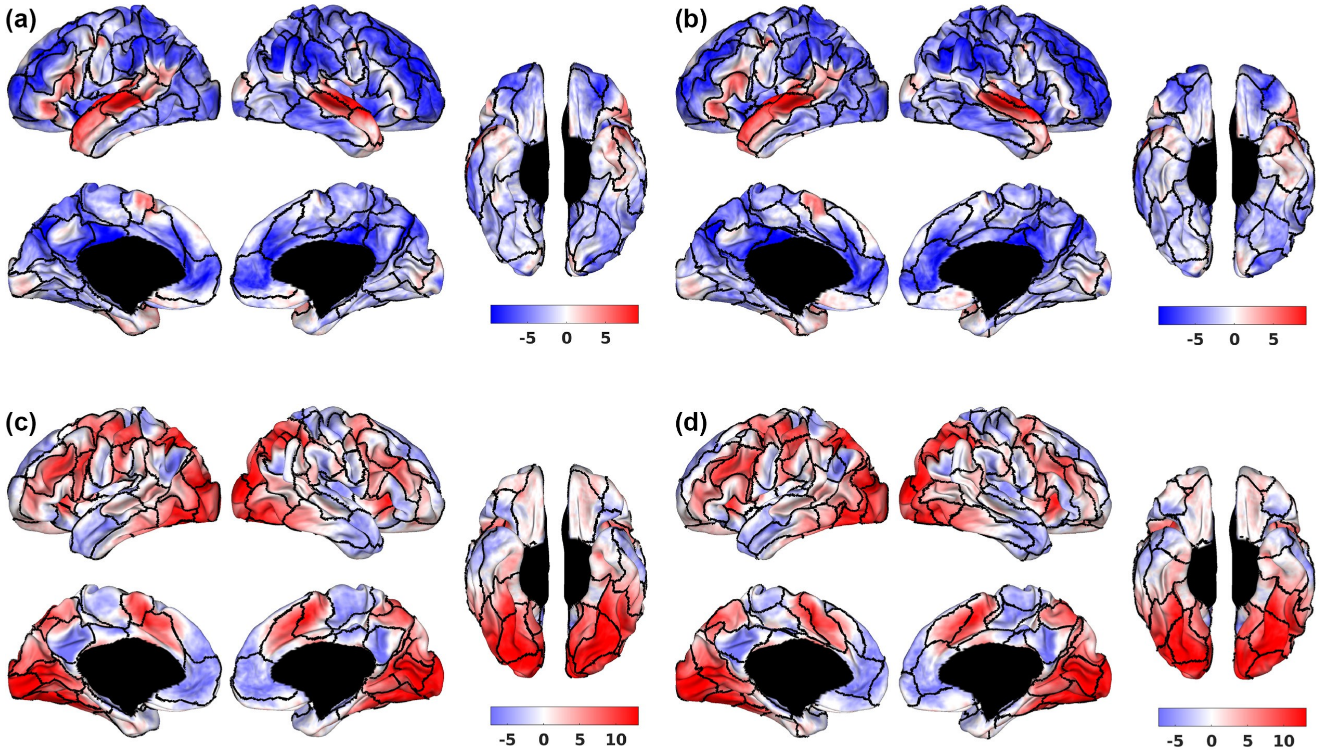

FIGURE 4 Example HCP group average task activation *z*-score maps for Schaefer-100 and Untamed-100. First row: *language_story* contrast overlaid on boundaries (black) of (a) Schaefer-100 and (b) Untamed-100. Second row: *relational_match* contrast overlaid on (c) Schaefer-100 and (d) Untamed-100.

3.5 | Ablation study of graph construction methods

We also explored the effect of graph construction as described in Section 2.10 by comparing results using the NASCAR-based adjacency matrix with that computed using the correlation of the RSFC matrix. All other aspects of processing were identical. Supplementary Fig. S2 shows the bar plots of RSFC weighted average homogeneity of NASCAR-based and RSFC-based methods. Evidently, the parcellations generated from the NASCAR-based adjacency matrix substantially outperformed those generated from Pearson-based adjacency in all cases ($p < .001$). This demonstrated an advantage of using the results of NASCAR tensor decomposition to identify spatial networks over directly using the correlation of RSFC.

3.6 | Is there an optimal number of parcels?

Fig. 6 illustrates the ratio of RSFC homogeneity and task contrast variance between the Untamed atlas and random parcellations for the HCP task dataset. For RSFC homogeneity, a larger ratio indicates comparatively higher homogeneity for the Untamed atlas relative to its random parcellation counterpart. Conversely, for task contrast variance, a smaller ratio indicates better performance by the Untamed atlas.

Both curves exhibit a similar trend: the advantage of the Untamed atlas over random parcellations increases initially, peaks, and then diminishes as the number of parcels continues to grow. The optimal number of parcels varies between the two modalities. Fig. 6 (a) indicates that the greatest relative advantage of the Untamed atlas over random parcellations occurs at fewer than 50 parcels per hemisphere for the homogeneity metrics as evidenced by the peak. The contrast variance ratio in Fig. 6 (b) remains relatively flat and close to its minimum from approximately 50-150 parcels with the advantage relative to a random parcellation diminishing approximately monotonically above 200 parcels.

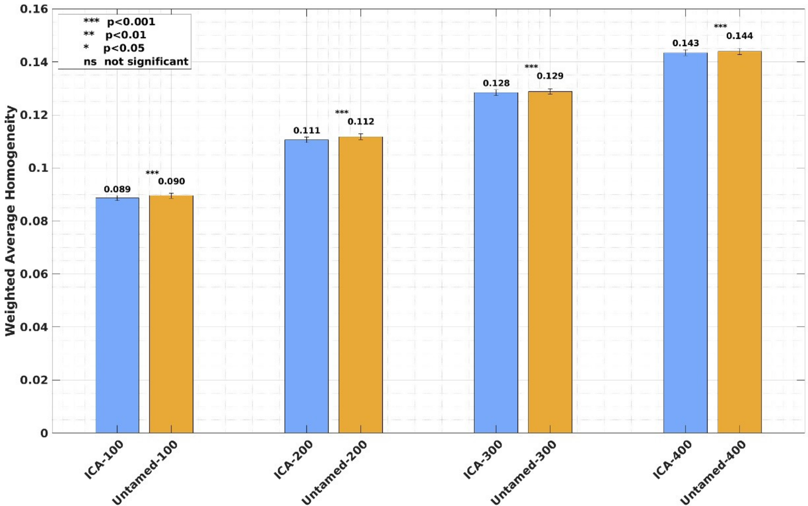

FIGURE 5 Weighted average RSFC homogeneity on the HCP dataset. Each bar plot depicts the subject-wise RSFC weighted average homogeneity, averaged across test subjects, for each atlas generated from ICA maps and the Untamed atlas (generated from NASCAR maps). Parcel numbers matched for the left and right hemispheres. The error bars represent the standard error across all subjects. The ICA baselines, derived from and optimized through hyperparameter tuning on the same HCP dataset, benefit from this dataset-specific tuning. In contrast, the Untamed atlas was generated using the independent GSP dataset with hyperparameters tuned on GSP. The ICA results represent the maximum homogeneity achievable under optimal hyperparameter settings for the HCP dataset.

3.7 | RSFC-based parcellation characterization

3.7.1 | Parcel-wise RSFC analysis

As noted in Section 3.1, the weighted-average RSFC homogeneity across the HCP subjects revealed no statistically significant differences between Untamed-300 and Schaefer-300, nor between Untamed-300 and Yan-300. The large sample size of 1000 test subjects indicated that the differences in homogeneity between Untamed and these two atlases at this particular parcel resolution were minimal. Nonetheless, a more detailed examination of the per-parcel RSFC homogeneity revealed nuanced localized differences. We conducted an in-depth comparison between Untamed-300 and Schaefer-300, summarized in Fig. 7. Figure 7(a) and (b) visualize the spatial distribution of parcel-wise RSFC homogeneity for Untamed-300 and Schaefer-300. Warmer colors indicate parcels with higher homogeneity values. These cortical maps show that both atlases produce broadly similar spatial patterns, although certain regions in association cortex exhibit localized differences in homogeneity. The spatial structure of these differences suggests that the modest deviations between the two parcellations are not uniformly random across the cortex. Figure 7(c) provides a quantitative view of these parcel-wise variations. The x-axis indexes the parcel number from 1 to 300. The left y-axis represents the RSFC homogeneity values averaged across HCP subjects for Untamed-300 (blue) and Schaefer-300 (orange). The right y-axis represents the Dice coefficient between corresponding parcels from the two atlases, which quantifies spatial overlap. This panel shows that although the global weighted-average homogeneity is similar, parcel-level values fluctuate across the cortex. Some parcels favor Untamed while others favor Schaefer. At the same time, the Dice coefficients remain high for most parcels, indicating that the overall spatial correspondence between the two atlases is strong. For completeness, Supplementary Fig. S3 and Fig. S4 show analogous results for Untamed-100 versus Schaefer-100 and Untamed-100 versus Random-100. For a direct comparison of spatial correspondence across multiple resolutions, Supplementary Fig. S5 shows Hungarian matching results between Untamed and Schaefer at 100, 200, 300, and 400 parcels, demonstrating consistent correspondence between the two atlas families across scales. The random parcellation, as described in Section 2.11, displays far weaker spatial alignment and more pronounced fluctuations in parcel-wise RSFC homogeneity.

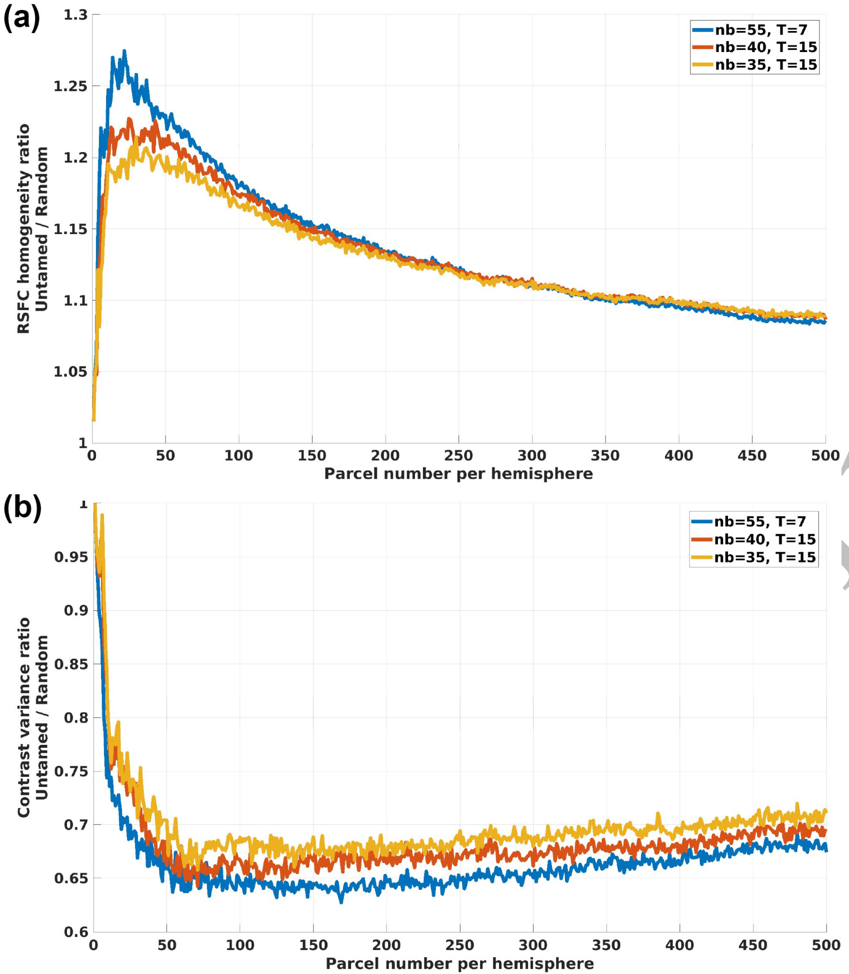

FIGURE 6 Ratios of three evaluation metrics comparing Untamed and random parcellations across 1 to 500 parcels per hemisphere: (a) Ratio of weighted average RSFC homogeneity: a comparison between Untamed parcellations and the mean values from 50 random parcellation trials (b) Ratio of weighted average task contrast variance: a comparison between Untamed parcellations and the mean values from 50 random parcellation trials.

3.7.2 | Network-level RSFC analysis

The ordered RSFC matrices and the corresponding 7-network partitions for both atlases are shown in Fig. 8. The spatial boundaries of the resulting networks follow similar trends for Untamed-200 and Schaefer-200, as evident in the RSFC matrices (Fig. 8(a)(b)).

Quantitatively, the networks derived from Untamed-200 achieved an RSFC homogeneity of 0.2077 ± 0.0143, slightly higher than the corresponding value for Schaefer-200 (0.2042 ± 0.0180, $p = 0.0696$, with paired sample t-test) across the Yale rsfMRI subjects. To assess alignment with canonical organization, network labels were projected back to the cortical surface (Fig. 8(c)(d)), and the Adjusted Rand Index (ARI; Section 2.14) was computed relative to Yeo's 7-network parcellation (Yeo et al., 2011). Untamed-200 achieved an ARI of 0.3742, compared to 0.3605 for Schaefer-200, indicating slightly stronger overall correspondence with Yeo-7 networks.

This global pattern, however, varies across individual networks. For instance, the visual and dorsal sensorimotor networks in Schaefer-200 show closer similarity to Yeo-7 networks, whereas the frontoparietal and ventral sensorimotor networks in Untamed-200 more closely match the Yeo-7 networks. These results highlight nuanced differences in network delineation between the two atlases, with each demonstrating strengths in different functional systems.

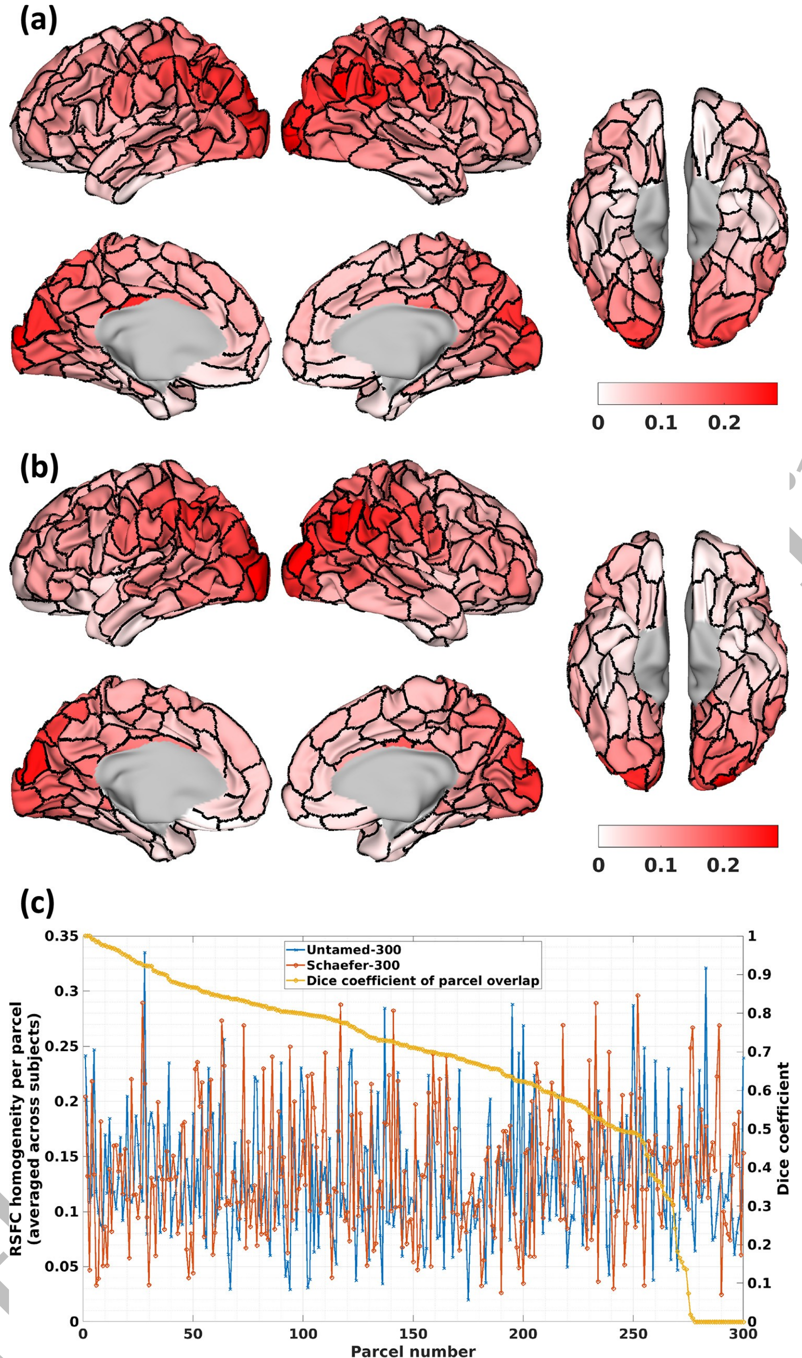

FIGURE 7 Parcel-wise RSFC homogeneity scores (averaged across the rsfMRI data of 1000 HCP subjects) visualized on the parcel boundaries for (a) Untamed-300 and (b) Schaefer-300 (c) Comparison of parcel-wise homogeneity scores between the two atlases, with parcels matched using the Hungarian matching algorithm. The rank-ordered Dice coefficients between matched parcels are also shown.

3.8 | NASCAR-derived network participation similarity

Revisiting the NASCAR networks that inform the Untamed parcels, Fig. 9 shows the 100- and 300-parcel versions of Untamed superimposed on two default mode subnetworks, as well as the visual and sensorimotor networks derived from NASCAR. The parcel boundaries of Untamed generally align with transitions between regions of higher and lower network expression in the

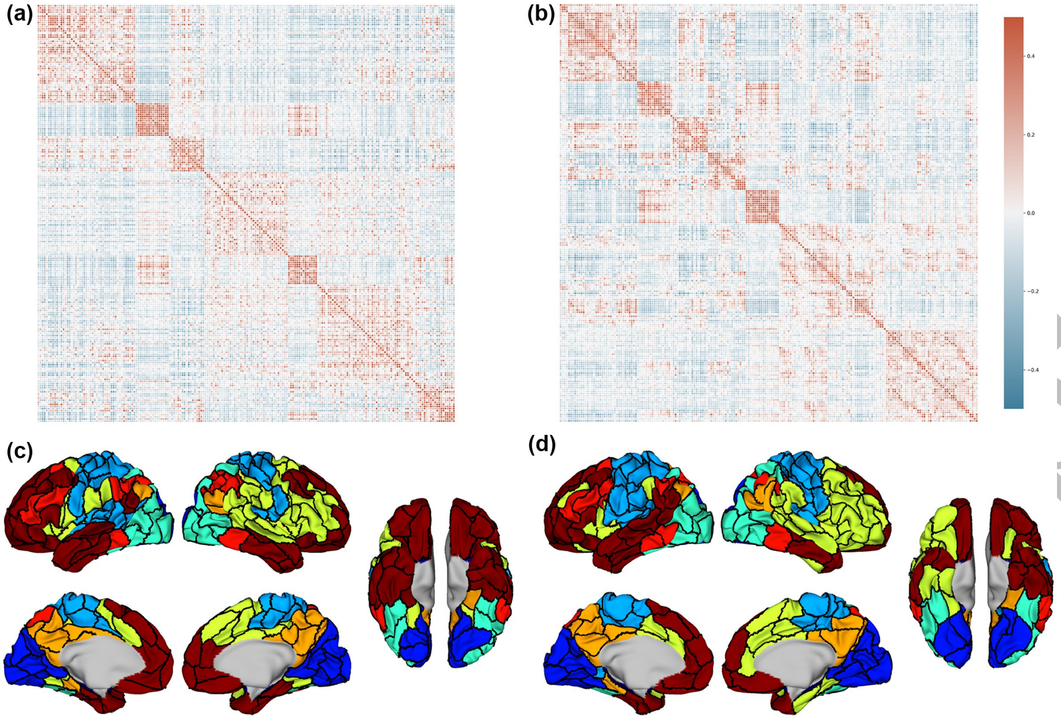

FIGURE 8 RSFC based on (a) Untamed-200 and (b) Schaefer-200 (averaged across subjects) on Yale rsfMRI test subjects after global signal regression. Both clustered to 7 networks as in (Yeo et al., 2011) using spectral clustering. (c) and (d) shows all parcels of Untamed-200 and Shaefer-200 assigned network colors.

NASCAR spatial maps, indicating that these functional gradients help guide the resulting parcellations. Additionally, Fig. 9 illustrates that a cortical vertex can exhibit nonzero loadings on multiple networks, reflecting the overlapping nature of the NASCAR components.

To further characterize this overlap, we examined vertex-wise similarity in network participation using the degree measure described in Section 2.13. The resulting degree distribution, shown in Fig. 10, highlights substantial variability across the cortex. Elevated degree values are concentrated in prefrontal, posterior cingulate, temporal, and parietal regions, which correspond to functional hubs reported in prior work (Buckner et al., 2009; Van Den Heuvel & Sporns, 2013). These findings underscore the ability of the overlapping and correlated NASCAR spatial maps to capture meaningful patterns of large-scale network participation.

3.9 | Reproducibility analysis

Reproducibility measured by ARI (Section 2.14) is shown in Fig. 11. The ARI between the two independently generated Untamed atlases (Group 1 vs. Group 2) and between each subgroup atlas and the full-sample atlas are the highest across all comparisons, indicating strong reproducibility of the Untamed parcellation across subject groups and varying sample sizes. Comparisons with existing atlases show moderate agreement: Schaefer–Yan exhibits consistently higher ARI values than Untamed–Schaefer or Untamed–Yan, reflecting the methodological similarity and shared dataset underlying Schaefer and Yan. As expected, agreement with the Random atlas is lowest at every parcel resolution. Overall, the ARI trends in Fig. 11 highlight the reproducibility of Untamed and the relative similarity structure among the different atlas families.

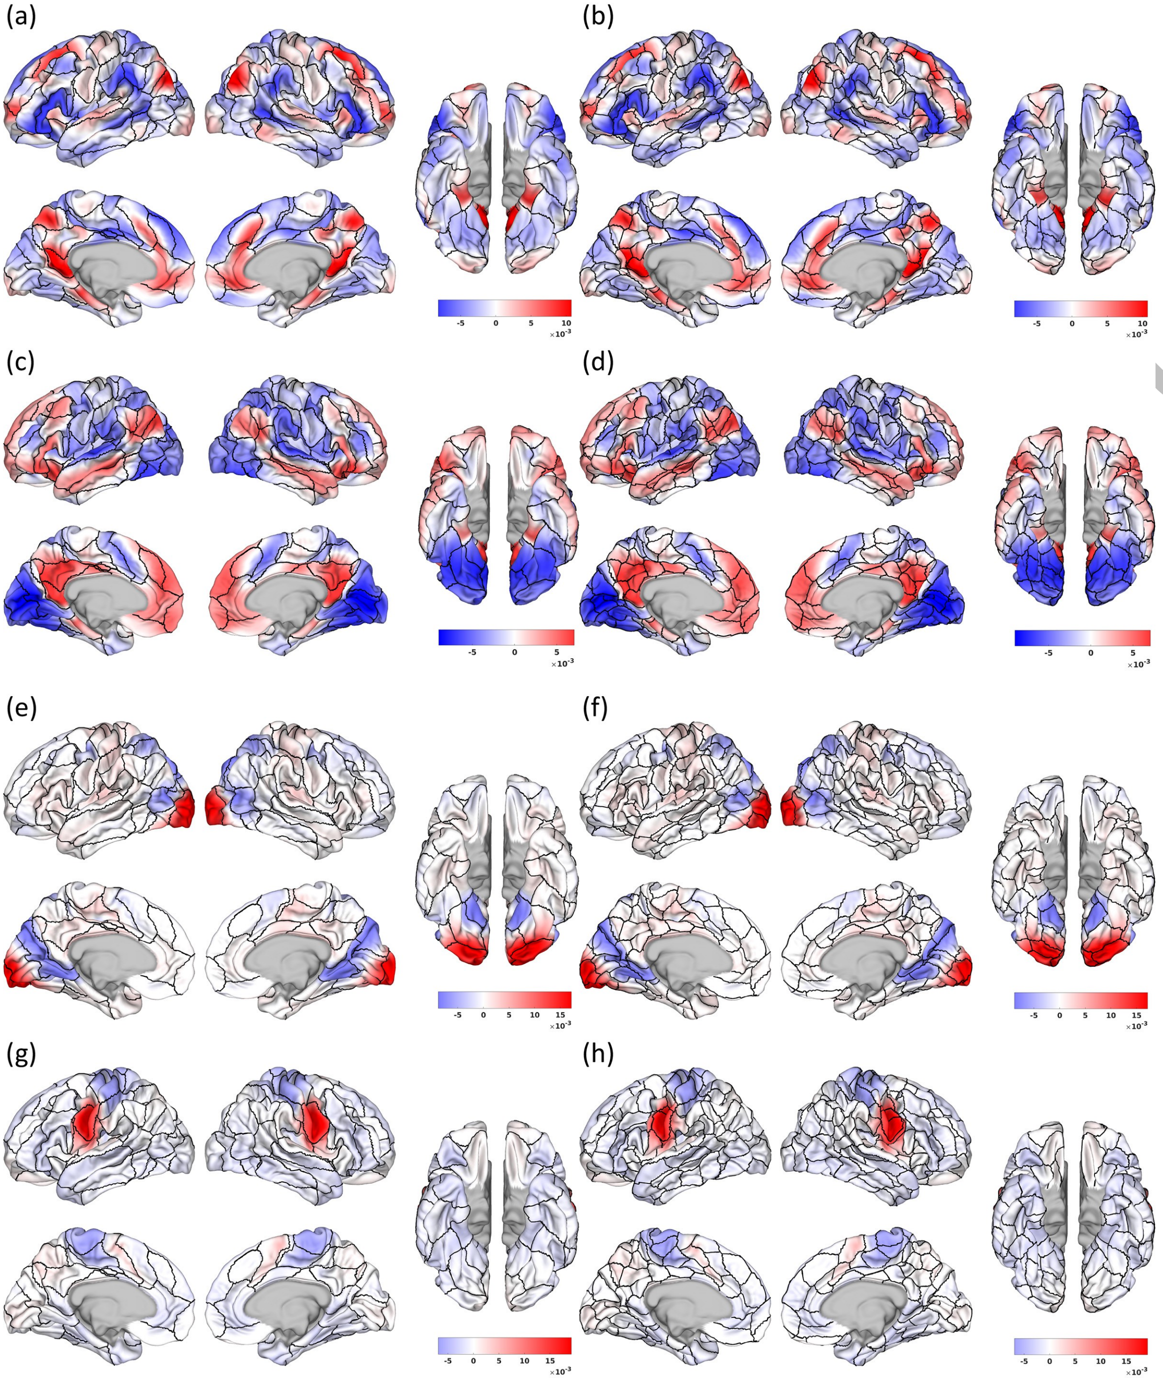

F I G U R E 9 (a)(c) and (b)(d) depict two NASCAR default mode sub-networks with (a)(c) Untamed-100 (b)(d) Untamed-300 parcel boundaries; (e) and (f) display a NASCAR visual network with (c) Untamed-100 and (d) Untamed-300; (g) and (h) display NASCAR sensorimotor network with (g) Untamed-100 and (h) Untamed-300.

4 | DISCUSSION

We introduced *Untamed*, a novel cortical parcellation scheme developed from population resting-state fMRI data. Untamed constructed spatially disjoint parcels by leveraging the correlated brain networks identified by NASCAR tensor decomposition (Li et al., 2021, 2023). Untamed generates spatially disjoint parcels whose boundaries follow spatial variations in network expression, and the resulting atlases exhibit strong correspondence with both resting-state connectivity patterns and task-evoked

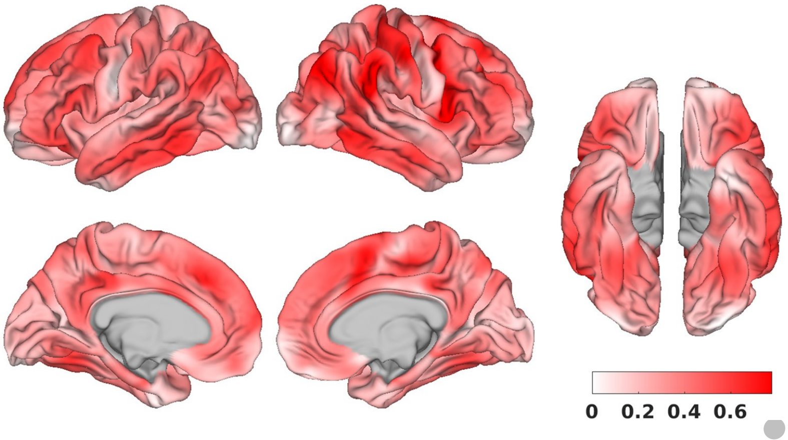

FIGURE 10 Degree distribution across the cortex, calculated from the spatial correlation of the NASCAR spatial maps and normalized to range within 0 and 1.

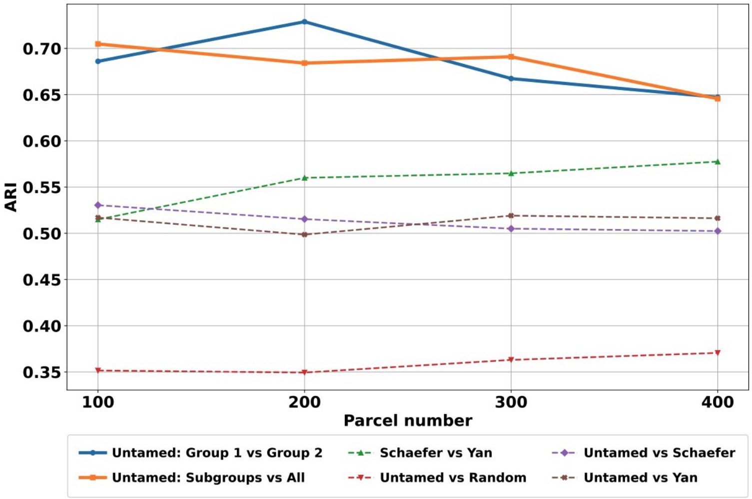

FIGURE 11 Adjusted Rand Index (ARI) across parcel numbers. The full set of GSP subjects was randomly split into two equal subgroups (Group 1 and Group 2, each with 714 subjects). The figure presents ARI scores between: (1) Untamed Group 1 v.s. Group 2, (2) Untamed subgroups v.s. All, (3) Schaefer v.s. Yan, (4) Untamed v.s. Random, (5) Untamed v.s. Schaefer, and (6) Untamed v.s. Yan.

functional organization. The results presented earlier demonstrate that Untamed performs competitively across a wide range of metrics, datasets, and spatial resolutions.

Beyond performance, our findings offer several insights for the broader parcellation and network neuroscience literature. First, the success of Untamed highlights the value of using network representations that do not impose statistical constraints such as independence or orthogonality. A growing body of work has emphasized that large-scale functional systems are not strictly segregated (Buckner et al., 2009; Yeo et al., 2011; Van Den Heuvel & Sporns, 2013), and that shared patterns of participation across networks are common. Methods that explicitly enforce independence, such as ICA, or rely on strong probabilistic priors on spatial smoothness and covariance structure, such as probabilistic functional mode approaches (Harrison et al., 2015), may therefore limit the expression of correlated network structure. In contrast, the NASCAR-derived spatial maps capture correlated patterns of network expression with minimal imposed constraints, providing a more flexible and biologically plausible representation of functional organization. As illustrated in Supplementary Fig. S6, NASCAR components exhibit rich correlation structure, whereas ICA components are tightly clustered around zero, reflecting the impact of independence constraints. By combining these unconstrained network representations with a graph-based embedding and clustering pipeline, Untamed translates continuous network structure into discrete, interpretable cortical parcels.

Our comparisons across multiple datasets and task domains illustrate how atlas choice can influence downstream analyses. Indeed, prior works have shown that atlas selection affects estimates of functional connectivity, topographic alignment, and structure-function correspondence (Bryce et al., 2021; Moghimi et al., 2022). By demonstrating that high-quality parcels can be derived systematically from the underlying network structure, Untamed suggests a shift toward more principled and data-informed parcellation strategies that reduce reliance on heuristic or convenience-based atlas choices.

The Untamed framework provides new opportunities for understanding the relationship between mesoscale parcellations and large-scale functional networks. Our network-level analyses indicate that parcels derived from NASCAR-based features can be grouped into coherent functional systems that align with canonical network templates while preserving finer variations across networks. This may facilitate studies that link localized functional units with distributed network organization, an important goal in current models of hierarchical brain function.

The practical advantages of Untamed also merit emphasis. The pipeline is a fully automated parcellation pipeline that requires minimal manual intervention and is not sensitive to the hyperparameter choices, which enhances its robustness and practical usability. Once the NASCAR-derived network features are obtained, atlases at any desired spatial resolution can be generated efficiently without reprocessing the underlying time series data. This flexibility allows researchers to tailor parcellation granularity to specific analytic goals while maintaining methodological consistency across resolutions.

A further insight from this work concerns the long-standing question of whether a natural or optimal number of cortical parcels can be identified from resting-state data. Prior studies have reached mixed conclusions, with some suggesting a hierarchy of scales rather than a single optimal resolution (Schaefer et al., 2018; Eickhoff et al., 2018). Our comparison against random parcellations indicates that the relative advantage of a principled atlas is not uniform across resolutions. Instead, the benefit of Untamed increases initially, reaches a peak, and gradually diminishes as the number of parcels becomes very large. This pattern supports the view that functional organization is neither scale-free nor concentrated at a single dominant resolution. Instead, a meaningful structure exists within a bounded range of scales, beyond which parcels become too fine to capture coherent functional units. These observations reinforce the idea that multiscale atlases and scale-aware analytic approaches may be more appropriate than attempting to identify a single optimal parcellation size.

Another advantage of our framework is that Untamed is compatible with both volume-space and surface-space representations, as the tensor decomposition, graph construction, embedding, and clustering steps do not depend on surface geometry. This enables broader applicability across diverse datasets and imaging modalities. The ability to integrate correlated network features with graph-based embedding and clustering also makes the methodology adaptable for specialized applications, such as parcellations tailored to clinical cohorts, multimodal datasets, or developmental populations.

One of the limitations of the current work is that our evaluations centered on young adults, largely due to the availability of benchmark datasets with well-established reference atlases. Extending Untamed to lifespan cohorts such as HCP-Development and HCP-Aging may reveal developmental or age-related shifts in functional topography and could test whether the overlapping network features that guide the method remain stable across the lifespan.

Additionally, while this study centered on group-level atlases to support direct comparison with widely used parcellations, the modular components of Untamed naturally support individualized parcellation. Approaches that integrate group priors with subject-specific refinements (Chong et al., 2017) represent a promising direction for generating personalized functional maps. Further, our decision to retain both hemispheres reflects the presence of subtle but meaningful functional asymmetries, including those related to language and attention (Gotts et al., 2013). However, a single-hemisphere variant may offer computational advantages and could be explored in future work, provided its impact on asymmetric functional organization is carefully evaluated.

To conclude, we introduced *Untamed* - a practical tool for functional parcellation and a platform for exploring the principles that shape cortical organization. By integrating overlapping network representations with a coherent parcellation framework, this work encourages a shift toward more flexible and biologically grounded approaches to defining the functional architecture of the human brain.

5 | DATA AND CODE AVAILABILITY

The data used in this study are publicly available from the Genomics Superstruct Project (GSP) (https://www.neuroinfo.org/gsp) and the Human Connectome Project, Young Adult Study (https://www.humanconnectome.org/study/hcp-young-adult). The MDTB dataset task contrast maps are publicly available at https://github.com/DiedrichsenLab/DCBC. The parcellations generated in this study and the associated code are available at the https://untamed-atlas.github.io.

References

Akrami, H., Joshi, A. A., Li, J., & Leahy, R. M. (2019). Group-wise alignment of resting fmri in space and time. In E. D. Angelini & B. A. Landman (Eds.), *Spie medical imaging 2019: Image processing* (p. 103). San Diego, United States. doi: 10.1117/12.2512564

Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T., & Calhoun, V. D. (2014). Tracking whole-brain connectivity dynamics in the resting state. *Cerebral Cortex*, *24*, 663–676. doi: 10.1093/cercor/bhs352

Andrews-Hanna, J. R., Reidler, J. S., Sepulcre, J., Poulin, R., & Buckner, R. L. (2010). Functional-anatomic fractionation of the brain's default network. *Neuron*, *65*, 550–562. doi: 10.1016/j.neuron.2010.02.005

Andrews-Hanna, J. R., Smallwood, J., & Spreng, R. N. (2014). The default network and self-generated thought: component processes, dynamic control, and clinical relevance. *Annals of the New York Academy of Sciences*, *1316*(1), 29–52. doi: 10.1111/nyas.12360

Arslan, S., Ktena, S. I., Makropoulos, A., Robinson, E. C., Rueckert, D., & Parisot, S. (2018). Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex. *NeuroImage*, *170*, 5–30.

Barch, D. M., Burgess, G. C., Harms, M. P., Petersen, S. E., Schlaggar, B. L., Corbetta, M., . . . Essen, D. C. V. (2013). Function in the human connectome: Task-fmri and individual differences in behavior. *NeuroImage*, *80*, 169–189. doi: 10.1016/j.neuroimage.2013.05.033

Beckmann, C. F., DeLuca, M., Devlin, J. T., & Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B*, *360*(1457), 1001–1013. doi: 10.1098/rstb.2005.1634

Biswal, B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, *34*, 537–541. doi: 10.1002/mrm.1910340409

Brodmann, K. (1909). *Vergleichende lokalisationslehre der großhirnrinde: in ihren prinzipien dargestellt auf grund des zellenbaues*. Leipzig: Barth.

Bryce, N. V., Flournoy, J. C., Guassi Moreira, J. F., Rosen, M. L., Sambook, K. A., Mair, P., & McLaughlin, K. A. (2021). Brain parcellation selection: An overlooked decision point with meaningful effects on individual differences in resting-state functional connectivity. *NeuroImage*, *243*, 118487. doi: 10.1016/j.neuroimage.2021.118487

Buckner, R. L., & DiNicola, L. M. (2019). The brain's default network: updated anatomy, physiology and evolving insights. *Nature Reviews Neuroscience*, *20*(10), 593–608. doi: 10.1038/s41583-019-0212-7

Buckner, R. L., Sepulcre, J., Talukdar, T., Krienen, F. M., Liu, H., Hedden, T., . . . Johnson, K. A. (2009). Cortical hubs revealed by intrinsic functional connectivity: Mapping, assessment of stability, and relation to alzheimer's disease. *Journal of Neuroscience*, *29*(6), 1860–1873. doi: 10.1523/JNEUROSCI.5062-08.2009

Chong, M., Bhushan, C., Joshi, A. A., Choi, S., Haldar, J. P., Shattuck, D. W., . . . Leahy, R. M. (2017). Individual parcellation of resting fmri with a group functional connectivity prior. *NeuroImage*, *156*, 87–100. doi: 10.1016/j.neuroimage.2017.04.054

Cichocki, A., Mandic, D., De Lathauwer, L., Zhou, G., Zhao, Q., Caiafa, C., & Phan, H. A. (2015). Tensor decompositions for signal processing applications: From two-way to multiway component analysis. *IEEE Signal Processing Magazine*, *32*(2), 145–163. doi: 10.1109/MSP.2013.2297439

Coalson, T. S., Van Essen, D. C., & Glasser, M. F. (2018). The impact of traditional neuroimaging methods on the spatial localization of cortical areas. *Proceedings of the National Academy of Sciences*, *115*(27), E6356–E6365.

Craddock, R. C., James, G. A., Holtzheimer, P. E., Hu, X. P., & Mayberg, H. S. (2012). A whole brain fmri atlas generated via spatially constrained spectral clustering. *Human Brain Mapping*, *33*(8), 1914–1928. doi: 10.1002/hbm.21333

Eickhoff, S. B., Yeo, B. T. T., & Genon, S. (2018). Imaging-based parcellations of the human brain. *Nature Reviews Neuroscience*, *19*(11), 672–686. doi: 10.1038/s41583-018-0071-7

Esteban, O., Markiewicz, C. J., Blair, R. W., Moodie, C. A., Isik, A. I., Erramuzpe, A., . . . Gorgolewski, K. J. (2019). fmriprep: a robust preprocessing pipeline for functional mri. *Nature Methods*, *16*(1), 111–116. doi: 10.1038/s41592-018-0235-4

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, *102*(27), 9673–9678.

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., . . . others (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, *536*(7615), 171–178.

Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., . . . Jenkinson, M.

(2013). The minimal preprocessing pipelines for the human connectome project. *NeuroImage*, *80*, 105–124. doi: 10.1016/j.neuroimage.2013.04.127

Gordon, E. M., Laumann, T. O., Adeyemo, B., Huckins, J. F., Kelley, W. M., & Petersen, S. E. (2016). Generation and evaluation of a cortical area parcellation from resting-state correlations. *Cerebral Cortex*, *26*(1), 288–303. doi: 10.1093/cercor/bhu239

Gotts, S. J., Jo, H. J., Wallace, G. L., Saad, Z. S., Cox, R. W., & Martin, A. (2013). Two distinct forms of functional lateralization in the human brain. *Proceedings of the National Academy of Sciences*, *110*(36), E3435–E3444.

Gramfort, A. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, *7*, 267. doi: 10.3389/fnins.2013.00267

Greve, D. N., & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, *48*(1), 63–72. doi: 10.1016/j.neuroimage.2009.06.060

Han, M., Yang, G., Li, H., Zhou, S., Xu, B., Jiang, J., . . . Gao, J.-H. (2020). Individualized cortical parcellation based on diffusion mri tractography. *Cerebral Cortex*, *30*(6), 3198–3208. doi: 10.1093/cercor/bhz303

Harris, K. D., & Mrsic-Flogel, T. D. (2013). Cortical connectivity and sensory coding. *Nature*, *503*(7474), 51–58. doi: 10.1038/nature12654

Harrison, S. J., Bijsterbosch, J. D., Segerdahl, A. R., Fitzgibbon, S. P., Farahibozorg, S.-R., Duff, E. P., . . . Woolrich, M. W. (2020). Modelling subject variability in the spatial and temporal characteristics of functional modes. *NeuroImage*, *222*, 117226. doi: 10.1016/j.neuroimage.2020.117226

Harrison, S. J., Woolrich, M. W., Robinson, E. C., Glasser, M. F., Beckmann, C. F., Jenkinson, M., & Smith, S. M. (2015). Large-scale probabilistic functional modes from resting state fMRI. *NeuroImage*, *109*, 217–231.

Holmes, A. J., Hollinshead, M. O., O'Keefe, T. M., Petrov, V. I., Fariello, G. R., Wald, L. L., . . . Buckner, R. L. (2015). Brain genomics superstruct project initial data release with structural, functional, and behavioral measures. *Scientific Data*, *2*, 150031. doi: 10.1038/sdata.2015.31

Hutchison, R. M., Womelsdorf, T., Allen, E. A., Bandettini, P. A., Calhoun, V. D., Corbetta, M., . . . others (2013). Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage*, *80*, 360–378.

Joshi, A. A., Choi, S., Liu, Y., Chong, M., Sonkar, G., Gonzalez-Martinez, J., . . . Leahy, R. M. (2022). A hybrid high-resolution anatomical mri atlas with sub-parcellation of cortical gyri using resting fmri. *Journal of Neuroscience Methods*, *374*, 109566. doi: 10.1016/j.jneumeth.2022.109566

Joshi, A. A., Chong, M., Li, J., Choi, S., & Leahy, R. M. (2018). Are you thinking what i'm thinking? synchronization of resting fmri time-series across subjects. *NeuroImage*, *172*, 740–752. doi: 10.1016/j.neuroimage.2018.01.058

Karahanoğlu, F. I., & Van De Ville, D. (2015). Transient brain activity disentangles fmri resting-state dynamics in terms of spatially and temporally overlapping networks. *Nature communications*, *6*(1), 7751.

King, M., Hernandez-Castillo, C. R., Poldrack, R. A., Ivry, R. B., & Diedrichsen, J. (2019). Functional boundaries in the human cerebellum revealed by a multi-domain task battery. *Nature Neuroscience*, *22*(8), 1371–1378. doi: 10.1038/s41593-019-0436-x

Kolda, T. G., & Bader, B. W. (2009). Tensor decompositions and applications. *SIAM Review*, *51*(3), 455–500. doi: 10.1137/07070111X

Kong, R., Li, J., Orban, C., Sabuncu, M. R., Liu, H., Schaefer, A., . . . Yeo, B. T. T. (2019). Spatial topography of individual-specific cortical networks predicts human cognition, personality, and emotion. *Cerebral Cortex*, *29*(6), 2533–2551. doi: 10.1093/cercor/bhy123

Lee, K., Horien, C., O'Connor, D., Garand-Sheridan, B., Tokoglu, F., Scheinost, D., . . . Constable, R. T. (2022). Arousal impacts distributed hubs modulating the integration of brain functional connectivity. *NeuroImage*, *258*, 119364. doi: 10.1016/j.neuroimage.2022.119364

Leech, R., Braga, R., & Sharp, D. J. (2012). Echoes of the brain within the posterior cingulate cortex. *Journal of Neuroscience*, *32*(1), 215–222.

Li, J., Kong, R., Liégeois, R., Orban, C., Tan, Y., Sun, N., . . . Yeo, B. T. T. (2019). Global signal regression strengthens association between resting-state functional connectivity and behavior. *NeuroImage*, *196*, 126–141. doi: 10.1016/j.neuroimage.2019.04.016

Li, J., Liu, Y., Wisnowski, J. L., & Leahy, R. M. (2023). Identification of overlapping and interacting networks reveals intrinsic spatiotemporal organization of the human brain. *NeuroImage*, *270*, 119944. doi: 10.1016/j.neuroimage.2023.119944

Li, J., Wisnowski, J. L., Joshi, A. A., & Leahy, R. M. (2021). Robust brain network identification from multi-subject asynchronous fmri data. *NeuroImage*, *227*, 117615. doi: 10.1016/j.neuroimage.2020.117615

Moghimi, P., Dang, A. T., Do, Q., Netoff, T. I., Lim, K. O., & Atluri, G. (2022). Evaluation of functional mri-based human brain parcellation: a review. *Journal of Neurophysiology*, *128*(1), 197–217. doi: 10.1152/jn.00411.2021

Ng, A., Jordan, M., & Weiss, Y. (2001). On spectral clustering: Analysis and an algorithm. In *Advances in neural information processing systems.*

Peng, L., Luo, Z., Zeng, L.-L., Hou, C., Shen, H., Zhou, Z., & Hu, D. (2023). Parcellating the human brain using resting-state dynamic functional connectivity. *Cerebral Cortex*, *33*(7), 3575–3590.

Perozzi, B., Al-Rfou, R., & Skiena, S. (2014). Deepwalk: online learning of social representations. In *Proceedings of the 20th acm sigkdd international conference on knowledge discovery and data mining* (pp. 701–710). New York, NY, USA: ACM. doi: 10.1145/2623330.2623732

Qiu, J., Dong, Y., Ma, H., Li, J., Wang, K., & Tang, J. (2018). Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *Proceedings of the eleventh acm international conference on web search and data mining* (pp. 459–467). Marina Del Rey, CA, USA: ACM. doi: 10.1145/3159652.3159706

Rockland, K. S. (2019). What do we know about laminar connectivity? *NeuroImage*, *197*, 772–784. doi: 10.1016/j.neuroimage.2017.07.032

Schaefer, A., Kong, R., Gordon, E. M., Laumann, T. O., Zuo, X.-N., Holmes, A. J., . . . Yeo, B. T. T. (2018). Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral Cortex*, *28*(9), 3095–3114. doi: 10.1093/cercor/bhx179

Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., . . . Greicius, M. D. (2007). Dissociable intrinsic connectivity networks for salience processing and executive control. *Journal of neuroscience*, *27*(9), 2349–2356.

Shen, X., Tokoglu, F., Papademetris, X., & Constable, R. T. (2013). Groupwise whole-brain parcellation from resting-state fmri data for network node identification. *NeuroImage*, *82*, 403–415. doi: 10.1016/j.neuroimage.2013.05.081

Smith, S. M., Hyvärinen, A., Varoquaux, G., Miller, K. L., & Beckmann, C. F. (2014). Group-PCA for very large fMRI datasets. *NeuroImage*, *101*, 738–749. doi: 10.1016/j.neuroimage.2014.07.051

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., . . . Joliot, M. (2002). Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain. *NeuroImage*, *15*(1), 273–289. doi: 10.1006/nimg.2001.0978

Van Den Heuvel, M. P., & Sporns, O. (2013). Network hubs in the human brain. *Trends in Cognitive Sciences*, *17*(12), 683–696. doi: 10.1016/j.tics.2013.09.012

Van Essen, D. C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T. E. J., Bucholz, R., . . . Yacoub, E. (2012). The human connectome project: A data acquisition perspective. *NeuroImage*, *62*(4), 2222–2231. doi: 10.1016/j.neuroimage.2012.02.018

Yan, X., Kong, R., Xue, A., Yang, Q., Orban, C., An, L., . . . Yeo, B. T. T. (2023). Homotopic local-global parcellation of the human cerebral cortex from resting-state functional connectivity. *NeuroImage*, *120010*. doi: 10.1016/j.neuroimage.2023.120010

Yeo, B. T. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., . . . Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*(3), 1125–1165. doi: 10.1152/jn.00338.2011

Zhi, D., King, M., Hernandez-Castillo, C. R., & Diedrichsen, J. (2022). Evaluating brain parcellations using the distance-controlled boundary coefficient. *Human Brain Mapping*, *43*(11), 3706–3720. doi: 10.1002/hbm.25878

Supplementary Information

Name-number of parcels	**Yale rsfMRI (Cohen's d)**	**HCP rsfMRI (Cohen's d)**
Yeo-51	0.2149	0.2117
Schaefer-100	0.0257	0.0176
Yan-100	0.0782	0.0826
Yeo-114	0.2372	0.1509
USCBrain-130	0.4494	0.3424
Power-130	0.6756	0.5342
Schaefer-200	0.0444	0.0190
Yan-200	0.0618	0.0288
Shen-200	0.5244	0.4241
Schaefer-300	0.0328	0.0001
Yan-300	0.0222	0.0004
Glasser-360	0.4058	0.3032
Schaefer-400	0.0239	-0.0151
Yan-400	0.0160	-0.0085

Table S1: Effect sizes (measured by Cohen's d) subject-wise RSFC homogeneity differences between Untamed and baseline atlases on Yale and HCP datasets. A positive Cohen's d indicates that Untamed introduces a higher RSFC homogeneity than the baseline counterpart, and vice versa.

nb	T	d	parcel-100	parcel-200	parcel-300	parcel-400
5	1	128	0.3878	0.5454	0.6441	0.7236
5	1	256	0.3838	0.5434	0.6420	0.7258
5	7	128	0.4090	0.5614	0.6593	0.7358
5	7	256	0.4135	0.5587	0.6608	0.7338
10	1	128	0.4174	0.5540	0.6484	0.7219
10	1	256	0.4166	0.5563	0.6434	0.7175
10	7	128	0.4261	0.5662	0.6616	0.7370
10	7	256	0.4263	0.5682	0.6631	0.7370
15	1	128	0.4244	0.5487	0.6498	0.7244
15	1	256	0.4235	0.5489	0.6433	0.7219
15	7	128	0.4292	0.5674	0.6634	0.7373
15	7	256	0.4287	0.5672	0.6619	0.7347
20	1	128	0.4186	0.5573	0.6556	0.7293
20	1	256	0.4232	0.5611	0.6546	0.7283
20	7	128	0.4298	0.5689	0.6646	0.7382
20	7	256	0.4284	0.5691	0.6641	0.7382
25	7	128	0.4312	0.5682	0.6649	0.7411
25	7	256	0.4304	0.5688	0.6640	0.7401
25	11	128	0.4357	0.5693	0.6651	0.7416
25	11	256	0.4340	0.5698	0.6671	0.7374
25	15	128	0.4363	0.5731	0.6670	0.7388
30	7	128	0.4339	0.5685	0.6645	0.7382
30	7	256	0.4360	0.5699	0.6653	0.7376
30	11	128	0.4355	0.5723	0.6634	0.7402
30	11	256	0.4318	0.5699	0.6662	0.7392
30	15	128	0.4377	0.5702	0.6680	0.7400
30	15	256	0.4352	0.5688	0.6675	0.7405
35	7	128	0.4362	0.5730	0.6672	0.7383
35	7	256	0.4354	0.5731	0.6675	0.7387
35	11	128	0.4403	0.5730	0.6673	0.7403
35	15	128	0.4378	0.5702	0.6662	0.7409
35	15	256	0.4400	0.5728	0.6665	0.7401
40	7	128	0.4405	0.5749	0.6646	0.7394
40	7	256	0.4423	0.5739	0.6639	0.7395
40	11	128	0.4464	0.5737	0.6688	0.7410
40	15	128	0.4444	0.5773	0.6689	0.7389
40	15	256	0.4412	0.5758	0.6661	0.7382
45	7	128	0.4450	0.5764	0.6667	0.7360
45	7	256	0.4491	0.5781	0.6654	0.7377
45	11	128	0.4433	0.5761	0.6651	0.7349
45	15	128	0.4466	0.5772	0.6652	0.7361
45	15	256	0.4459	0.5745	0.6651	0.7372
50	7	128	0.4462	0.5766	0.6675	0.7370
50	7	256	0.4497	0.5769	0.6661	0.7367
50	11	128	0.4477	0.5770	0.6663	0.7343
50	11	256	0.4494	0.5763	0.6618	0.7353
50	15	128	0.4492	0.5759	0.6651	0.7331
50	15	256	0.4477	0.5762	0.6642	0.7367
55	7	128	0.4506	0.5776	0.6648	0.7337
55	7	256	0.4482	0.5766	0.6656	0.7349
55	11	128	0.4443	0.5744	0.6675	0.7346
55	11	256	0.4484	0.5762	0.6658	0.7350
55	15	128	0.4481	0.5777	0.6624	0.7332
55	15	256	0.4501	0.5755	0.6668	0.7352
60	7	128	0.4494	0.5756	0.6649	0.7324
60	11	128	0.4484	0.5783	0.6659	0.7305
60	15	128	0.4459	0.5780	0.6639	0.7319

Table S2: Parameter sweep results.

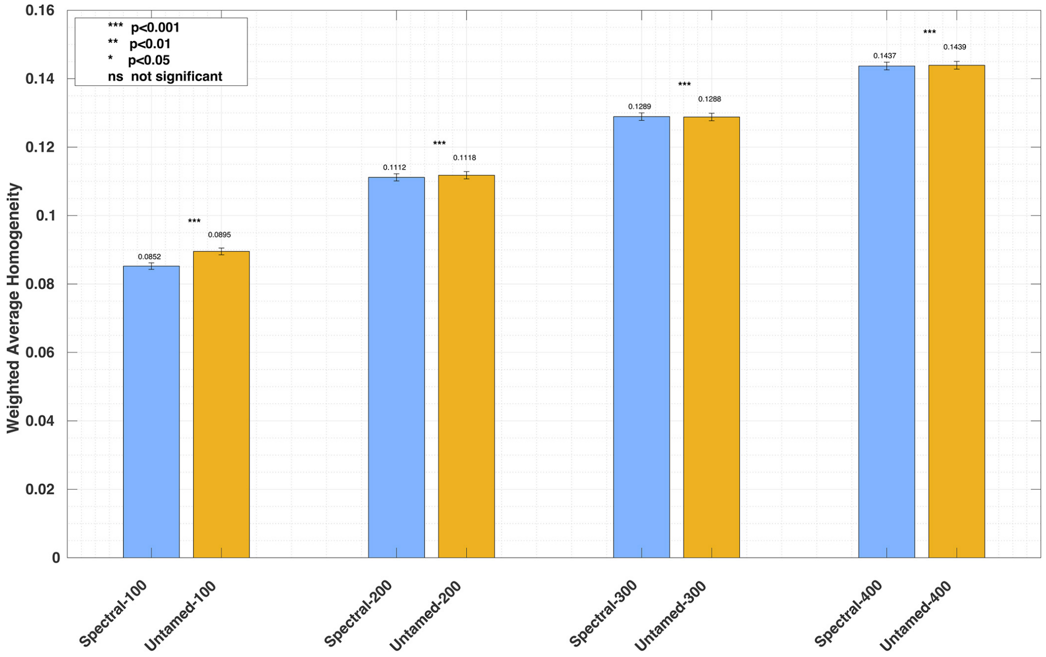

Figure S1: RSFC weighted average homogeneity across different cluster numbers matching existing parcellations. Each bar plot depicts the subject-wise RSFC weighted average homogeneity, averaged across 1000 HCP test subjects, for each spectral clustering-based atlas and the Untamed atlas, with matched parcel numbers for the left and right hemispheres. The error bars represent the standard error across all subjects. The comparison highlights atlases derived using NetMF-based features for clustering (Untamed) versus those generated via spectral clustering using Graph Laplacian eigenvectors.

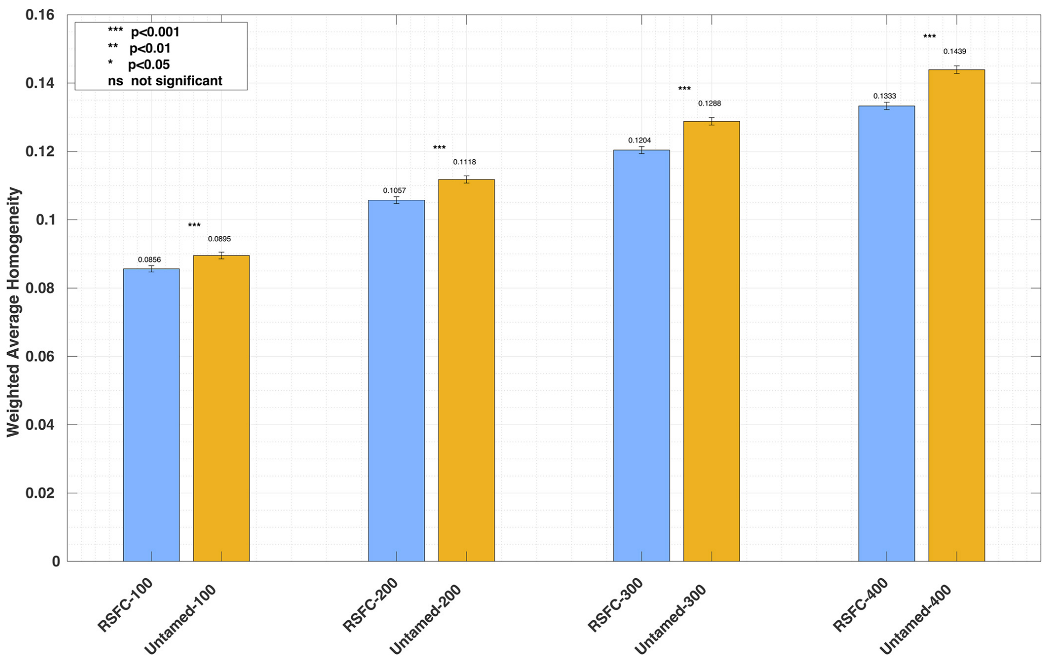

Figure S2: Weighted average RSFC homogeneity on the HCP dataset. Each bar plot depicts the subject-wise RSFC weighted average homogeneity, averaged across the 1000 test subjects, for each RSFC-based and the Untamed atlas, with matched parcel numbers for the left and right hemispheres. The error bars represent the standard error across all subjects. The RSFC-based average weighted homogeneity comparing NASCAR-based estimation of the graph adjacency matrix (used in Untamed) with computation from correlation of RSFC (the Pearson correlation of Pearson correlation between rsfMRI time-series).

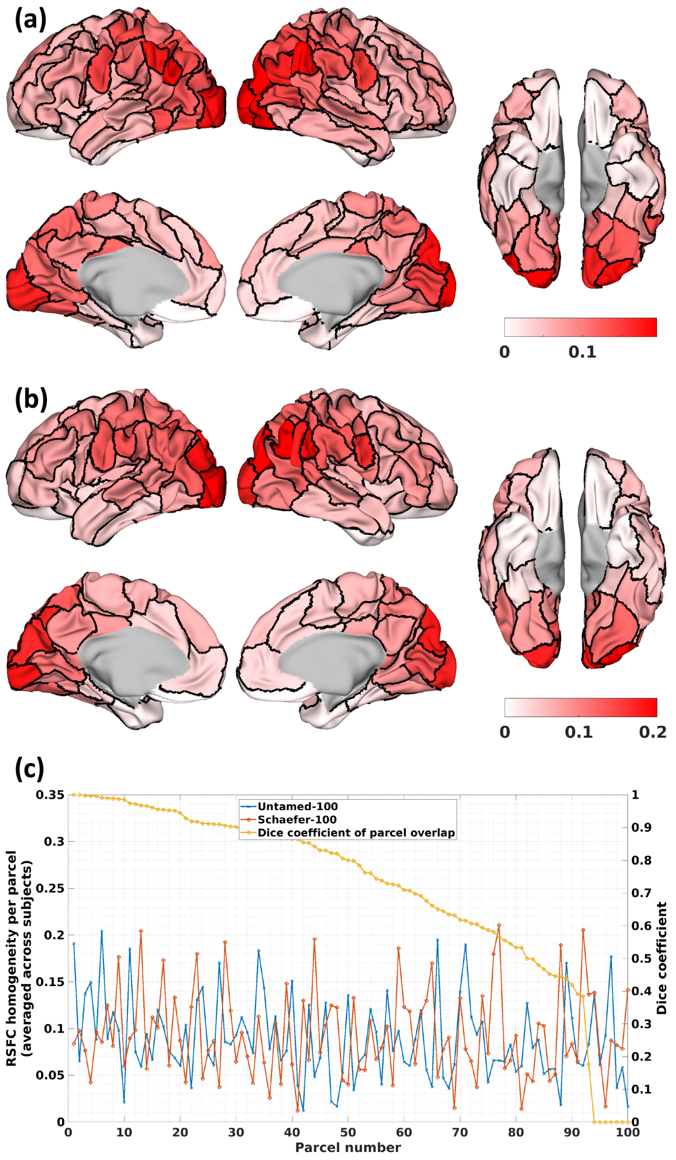

Figure S3: Parcel-wise RSFC homogeneity scores (averaged across the rsfMRI data of 1000 HCP subjects) visualized on the parcel boundaries for (a) Untamed-100, and (b) Schaefer-100. (c) Comparison of parcel-wise homogeneity scores between the two atlases, with parcels matched using the Hungarian matching algorithm. The rank-ordered Dice coefficients between matched parcels are also shown.

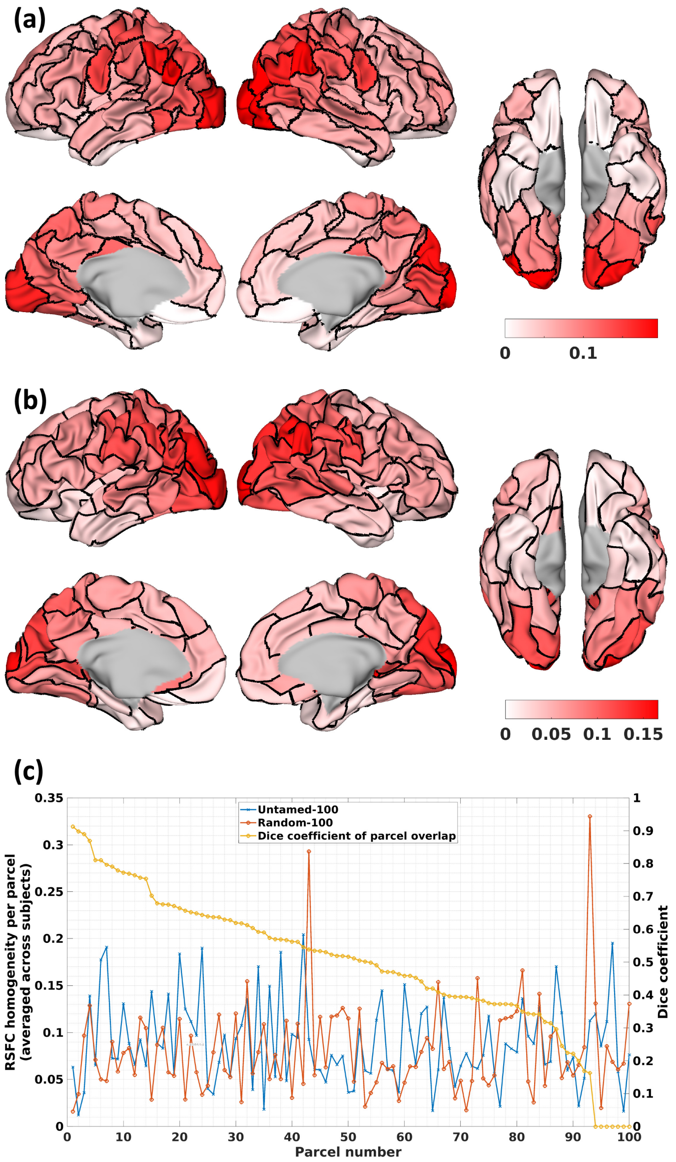

Figure S4: Parcel-wise RSFC homogeneity scores (averaged across the rsfMRI data of 1000 HCP subjects) visualized on the parcel boundaries for (a) Untamed-100 and (b) Random-100. (c) Comparison of parcel-wise homogeneity scores between the two atlases, with parcels matched using the Hungarian matching algorithm. The rank-ordered Dice coefficients between matched parcels are also shown.

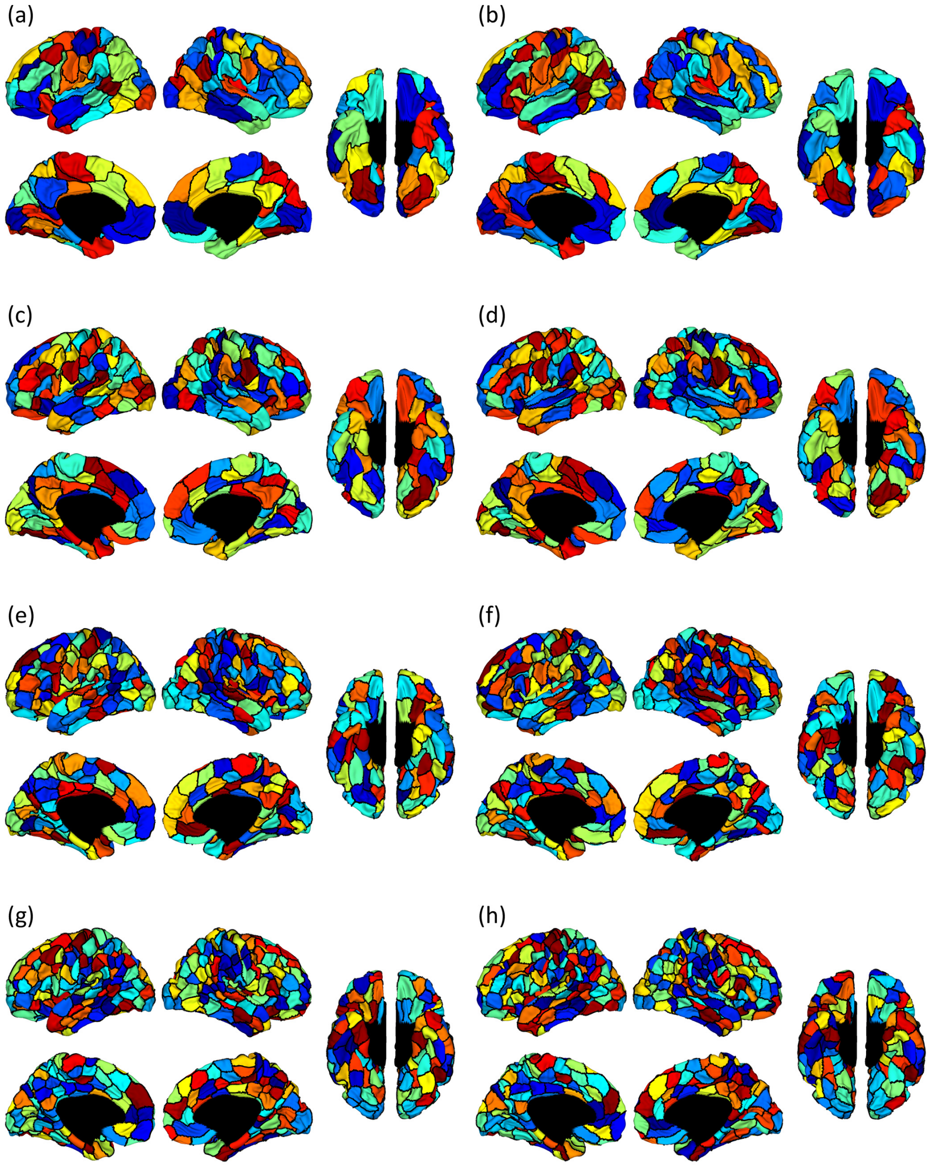

Figure S5: Parcellation comparison using Schaefers method: (a)(c)(e)(g): Schaefer-100/200/300/400 and Untamed: (b)(d)(f)(h): Untamed-100/200/300/400. From top to bottom we show 100, 200, 300, and 400 parcels. At each resolution, Hungarian matching was performed to match between the two parcellations to find maximum correspondence in terms of vertex-wise label assignment.

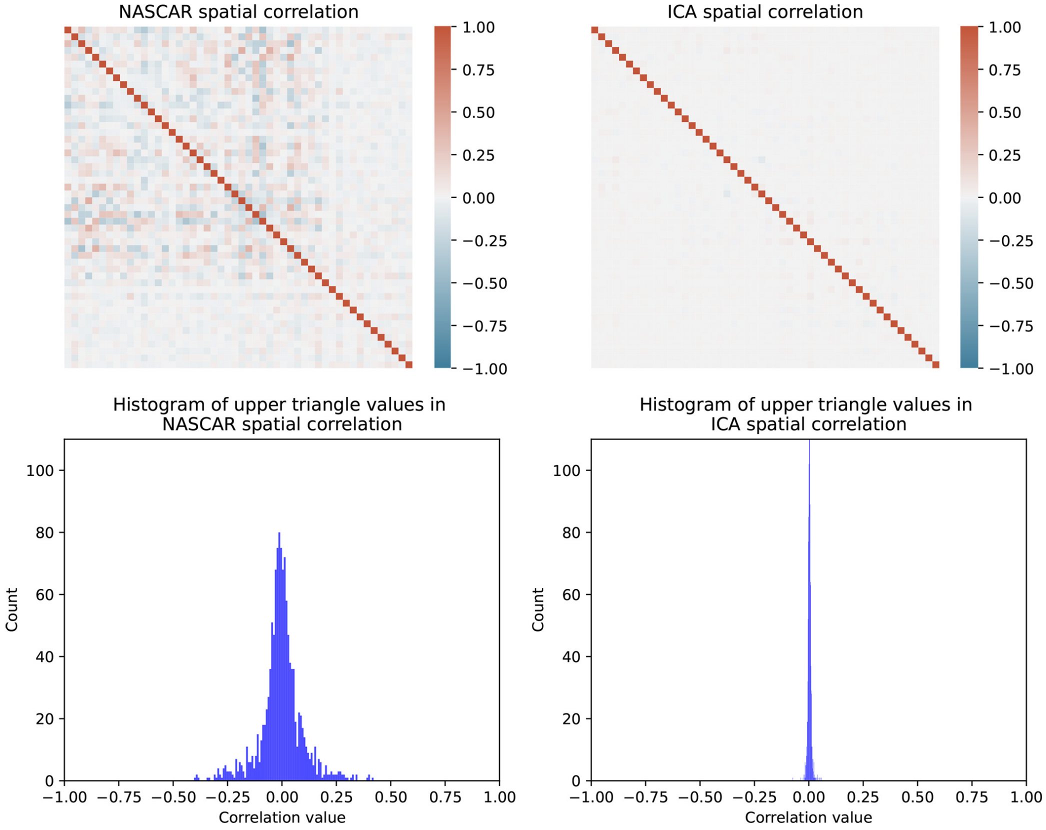

Figure S6: **First row:** Spatial correlation maps derived from pairwise Pearson correlations of the 50 spatial maps from NASCAR (left) and ICA (right). The ICA spatial maps were obtained from the HCP website (HCP 1200 project, computed from 1003 subjects). **Second row:** Histograms display the distribution of the upper triangular elements in each spatial correlation matrix. The ICA values are tightly clustered around zero, reflecting numerical noise. In contrast, the NASCAR values exhibit a broader distribution, indicating a lower degree of constraint imposed by the algorithm