Personalized whole-brain Ising models with heterogeneous nodes capture differences among brain regions

Adam Craig [1,2], Sida Chen[1,2], Qianyuan Tang[1,3], Changsong Zhou[1,2,3,4*]

[1]Department of Physics, Hong Kong Baptist University, Kowloon Tong, Hong Kong
[2] Centre for Nonlinear Studies and Beijing-Hong Kong-Singapore Joint Centre for Nonlinear and Complex Systems (Hong Kong), Hong Kong Baptist University, Kowloon Tong, Hong Kong
[3]Institute of Computational and Theoretical Studies, Hong Kong Baptist University, Kowloon Tong, Hong Kong
[4] Life Science Imaging Centre, Hong Kong Baptist University, Kowloon Tong, Hong Kong

*Corresponding author

E-mail: cszhou@hkbu.edu.hk
Department of Physics, Chair Professor
Life Science Imaging Centre, Director
Center for Nonlinear Studies, Director
Institute of Computational and Theoretical Studies, Director
Hong Kong Baptist University

Abstract

Multiple lines of research have studied how complex brain dynamics emerge from underlying connectivity by using Ising models as simplified neural mass models. However, limitations on parameter estimation have prevented their use with individual, high-resolution human neuroimaging data. Furthermore, most studies focus only on connectivity, ignoring node heterogeneity, even though real brain regions have different structural and dynamical properties.

Here we present an improved approach to fitting Ising models to 360-region functional MRI data: derivation of an initial guess model from group data, optimization of simulation temperature, and two stages of Boltzmann learning, first with group data, then with individual data. We then analyze how data binarization threshold affects goodness-of-fit, the role of the external field in model behavior, and correlations between model parameters and features from structural MRI, including myelination and sulcus depth.

The fitted models can reproduce the functional connectivity of the data across a wide range of thresholds, but the role of the external field parameters in model behavior increases with threshold. In parallel, the coupling between nodes correlates with structural connectivity throughout this range, but correlations between the external field and structural features increase with threshold. These results show how our methodology enables personalized, biophysically interpretable modeling of structure-function relationships at the whole-brain level, which can aid understanding of individual differences in brain network organization and dynamics. This approach will help to bridge the gap between connectomics, which emphasizes brain networks, and translational neuroscience, which often focuses on the unique roles of brain regions.

1. Introduction

1.1. The heterogeneous brain

Understanding how individual differences in brain structure shape differences in brain dynamics remains a central goal of modern neuroscience. The study of connectomics represents one prominent paradigm, in which researchers have made significant progress in linking individual differences in structural connectivity (SC), functional connectivity (FC), and cognitive abilities (Zimmermann et al., 2018; Wang et al., 2019b; Liégeois et al., 2020). However, the brain is not just a network of interchangeable nodes. Large-scale brain network analyses have revealed that microstructural features, such as myelination or cortical granularity, can significantly modulate local excitation-inhibition balance and structure-function coupling across brain regions (Glasser et al., 2014; Fotiadis et al., 2023; Fotiadis et al., 2024). Heterogeneity of local excitability can in turn influence larger-scale brain dynamics (Abeywardena, 2024). Importantly, individual differences in

the distribution of structural features across the brain can also predict differences in cognitive traits (Kristanto et al., 2023), suggesting that brain regions influence behavior through both their locations in the larger network and through their intrinsic properties. However, the high dimensionality of MRI data and the complexity of integrating single-subject functional and structural data make translating from correlations to mechanistic insights challenging (Bijsterbosch, et al., 2020; Venkadesh & Van Horn, 2021; Sultana et al., 2024).

1.2. The need for explainable, data-driven models

To meet this need, network-based approaches have shifted from static structure-function correlations to dynamical models with varying degrees of personalization (Lynn & Bassett, 2019; Guo et al., 2025). Such models could enable simulations of how the brain will respond to stimulation of different brain regions, leading to therapies tailored to the needs of each patient (Lynn & Bassett, 2019; Guo et al., 2025).

One extreme example of the detail now possible in dynamical models of the human brain is the Digital Twin Brain Project, which simulated 86 billion leaky integrate-and-fire neurons with 47.8 trillion synapses informed by multiple kinds of *in vivo* human data (Lu et al., 2024). However, it also illustrates the pitfalls of such a model, such as the risk of overfitting a high-dimensional model to limited data and the impracticality of comparing simulations of multiple individuals. Even before reaching this scale, the modeling community realized the need for approximations at macroscopic scales and developed multiple ways of modeling collective neural behavior, including neural mass models (Breakspear, 2017, Bansal et al., 2018). Both early neural mass models of animal nervous systems (Zhao et al., 2010) and newer models based on human brain data have suggested that regional heterogeneity and SC together shape patterns of FC (Demirtaş et al., 2019, Wang et al., 2019a, Deco et al., 2021).

However, these also face tradeoffs, as more complex models suffer from lower reliability despite greater individual specificity (Domhof et al., 2022). One key problem is that functional MRI (fMRI) provides an accurate view of spatial correlations but not temporal correlations in the underlying brain activity (Tu et al., 2024). This suggests that using a simpler model that captures spatial patterns of integration and segregation but is robust against spurious temporal fluctuations could provide a better bridge between brain structure and brain function.

1.3. Ising models as neural mass models

Over the past decade, the Ising model has proven capable of modeling many kinds of systems, because, when fitted to data, such models can capture the mean activity of and pairwise correlations between variables (Nguyen & Zecchina, 2017). After fitting, their simplicity makes their behavior explainable and their parameters amenable to a wide

variety of analyses (Nguyen & Zecchina, 2017). Most relevantly, a growing number of researchers have used all-to-all connected (generalized) Ising models as simplified neural mass models. When diffusion-MRI-derived SC serves as the network of couplings between regions, simulated FC correlates with fMRI-derived empirical FC, albeit modestly (correlation < 50%) (Das et al., 2014; Abeyasinghe et al., 2018). Despite the limitations of this approach, researchers have used it to model information transfer between nodes (Marinazzo et al., 2014), disorders of consciousness (Abeyasinghe et al., 2020), and the effects of anesthesia (Stramaglia et al., 2017), audio stimulation (Kandeepan et al., 2020), and sedatives (Kandeepan et al., 2020).

Other studies have achieved better correlations between model and data FC by fitting the model parameters directly to fMRI data, or, less commonly, to intracranial electroencephalogram data (Ashourvan et al., 2021). Notable downstream analyses of fMRI-fitted Ising models include clustering model states by the energy barriers between them (Ezaki & Watanabe, 2017), identifying important subnetworks through parameter sensitivity (Ponce-Alvarez et al., 2022), and calculating the correlation between the coupling parameters and SC (Ruffini et al., 2023). However, these examples pooled data from multiple individuals due to the data-intensive nature of the pseudolikelihood maximization fitting method they employed. This limitation to population-level models prevents the study of individual differences.

This inability to fit Ising models to individual data represents a major shortcoming compared to maximum likelihood gradient ascent, which other studies have used to fit models to single-subject fMRI data (Watanabe et al., 2013; Theis et al., 2024; Theis et al., 2025). These models have shown consistent and diagnostically meaningful individual differences in energy distribution among states (Theis et al., 2024; Theis et al., 2025). However, due to the need to compute the energies of all possible states, the authors only modeled small networks of fewer than 10 regions. Bayesian approximations allow fitting of larger 15-region models to individual data and show similar goodness of fit (Khanra et al., 2024). Parameters from individual models fitted with Bayesian approximations correlated significantly (around 50%-60%) with cognitive scores (Kang et al., 2021) or attention deficit and hyperactivity disorder-associated traits (Jeong et al., 2021). This apparent significance of individual differences in Ising model parameters suggested that being able to use a more fine-grained parcellation to identify areas of interest will improve the utility of Ising models in future practical applications. Specifically, many of the most promising targets for treating psychiatric disorders with transcranial magnetic stimulation are in higher-order areas, such as the prefrontal cortex and show high inter-patient variability in SC, making precise, personalized targeting valuable (Cash & Zalesky, 2024).

Choice of binarization threshold for the fMRI time series is an additional, insufficiently explored consideration. Recent work showed that threshold choice can reshape the

inferred energy landscape even in 10-node models (Theis et al., 2025), suggesting that it could have an even greater impact at whole-brain scales. Works that have addressed alternate choices of threshold have mainly considered them in terms of how they affected goodness of fit (Watanabe et al., 2013; Fortel et al., 2022; Theis et al., 2025), not how thresholding influenced the external fields, which are potentially important as the only distinguishing intrinsic features of nodes.

More generally, the role of the external field in the Ising model has received relatively little attention. Studies that use the SC for the coupling matrix typically drop the external field from the model entirely (Das et al., 2014; Marinazzo et al., 2014; Stramaglia et al., 2017; Abeyasinghe et al., 2018; Abeyasinghe et al., 2020; Fortel et al., 2022), though (Kandeepan et al., 2020) uses it to represent an external stimulus. No prior work involving Ising models fitted to fMRI data has specifically examined the importance of heterogeneity of external field to goodness of fit (Watanabe et al., 2013; Ezaki et al., 2017; Jeong et al., 2021; Kang et al., 2021; Ponce-Alvarez, 2022; Ruffini et al., 2023; Khanra et al., 2024; Theis et al., 2024; Theis et al., 2025). The studies that considered structure-function relationships only compared couplings and SC (Watanabe et al., 2013; Fortel et al., 2022; Ruffini et al., 2023). Of the two studies that use sparse canonical correlation analysis to compare individualized Ising model parameters to cognitive or behavioral scores, (Kang et al., 2021) uses the external field and couplings together, while (Jeong et al., 2021) uses only the couplings. The heterogeneity of nodes in the fitted Ising model represents an overlooked opportunity to connect structural features to local excitability and local excitability to whole-brain-scale dynamics.

1.4. Personalized whole-brain Ising models with heterogeneous nodes

Here we present a GPU-accelerated, four-stage workflow that fits fully connected 360-region Ising models to individual HCP resting-state fMRI data, which we then apply to an analysis of the influence of binarization threshold on goodness of fit, node heterogeneity, and correlations between external field and brain structural features. Our approach builds on Boltzmann learning with Metropolis Monte Carlo sampling, which others have used for fitting Ising models to spiking neuron recordings (Roudi et al., 2009). While computationally intensive, it offers a high degree of reliability when given sufficient sampling simulation length and parameter update steps (Roudi et al., 2009). We improved the reliability of the method and decreased the number of fitting steps needed by starting from a data-derived initialization and optimizing simulation temperature. We also made use of GPU acceleration to generate hundreds of personalized models in parallel.

To quantify how binarization threshold influences goodness of fit and the role of heterogeneous external fields, we first compared different thresholds applied to the same pooled group data. As in (Watanabe et al., 2013; Fortel et al., 2022; Theis et al., 2025), we found that fitting was fastest and most reliable when we binarized data at the mean, at

which threshold we also observed a correlation between model coupling and SC close to that reported by (Ruffini et al., 2023). However, we found that fitting succeeded at thresholds up to 1 SD above the mean and that this higher threshold was necessary to elicit meaningful node heterogeneity that correlated with region-to-region differences in myelination, curvature, and other cortical features measured with structural MRI. We then used this threshold to fit single-subject models and found that individual differences in external field contributed to capturing individual differences in FC. Finally, we quantified the local structure-function coupling of each region as the predictability of individual differences in Ising model external field from individual differences in structural features.

To the best of our knowledge, this is the first work to demonstrate fitting of 360-region whole-brain Ising models to single-subject data or to use Ising model external field parameters to study the relationship between local excitability and cortical structural features. In summary, our results provide a principled route toward structure-informed, subject-specific Ising models that provide insights into how local microstructure and global network topology jointly sculpt brain dynamics, an essential step toward mechanistic biomarkers and precision neuro-stimulation planning.

Figure 1. Ising model fitting and analysis workflow. (**a**) Preprocess fMRI data to get a continuous-value time series for each region. (**b**) Binarize the continuous-valued data (blue) using a threshold expressed in terms of standard deviations (red) above the mean (green) of the single region time series. (**c**) Optimize the model parameters, the external field of a region (h_i) and coupling between regions (J_{ij}), using Boltzmann learning, sampling states with the Metropolis

algorithm. (**d**) Run a longer test simulation of the model to find the functional connectivity (FC). Then find the correlation between model and binarized data FC. (**e**) Find correlations between the fitted external field parameter value h_i and region structural features (thickness, myelination, curvature, and sulcus depth) from T1- and T2-weighted structural MRI. (**f**) Find the correlation between the coupling parameter J_{ij} and structural connectivity from DTI data.

2. Methods

2.1. Human brain MRI data

The data to which we fit the Ising models is resting-state fMRI data of healthy young adults from the Human Connectome Project (HCP) S1200 data release (Elam et al., 2021). We selected 837 individuals who each had a complete set of imaging data, including four fMRI time series, each 1200 time points sampled 0.72 seconds apart (14.388 minutes per scan), a DTI scan, a T1-weighted scan, and a T2-weighted scan. See (WU-Minn. et al., 2017) for more information about data collection procedures.

We preprocessed the fMRI time series in the manner used in (Wang et al., 2021). This pipeline started with the HCP minimal preprocessing pipeline (Glasser et al., 2013), which performed basic cleaning and alignment. We then assigned voxels to regions in the 360-region parcellation from (Glasser et al., 2016) to produce a 360×1200 time series for each scan. Finally, we band-pass filtered each region time series to the infra-slow band (0.01-0.1 Hz), the frequency range in which the blood-oxygenation-level-dependent (BOLD) signal components correlate strongly with infra-slow changes in electrophysiological and behavioral changes (Palva & Palva, 2012). For fitting results with the same HCP data parcellated with the 117-region Automatic Anatomical Labeling Atlas (ALL) Version 1 (Tzourio-Mazoyer et al., 2002), see Supplementary Figure S1.

A key stage of creating the fitting targets was binarization of the data, which required selection of a threshold. As in (Theis et al., 2025), we defined the threshold in terms of standard deviations (SD) above the mean of the local region time series and mapped values below the threshold to -1 and values above it to +1. To create an individual fitting target, we concatenated the four binarized time series of individual scans into a single 360 × 4800 sequence and took the means ($\langle\sigma_i\rangle_{\text{data}}$) and uncentered covariances ($\langle\sigma_i\sigma_j\rangle_{\text{data}}$). To create a group target, we randomly selected a training set of 670 subjects and took the means and uncentered covariances of their concatenated binarized time series data. We set aside the remaining 167 individuals as the testing set to test whether using the group Ising model as the initial guess for the individualized Ising model fitting would lead to differences in goodness of fit between individuals in and out of the training set.

We used the probabilistic tract-tracing method from (Behrens et al., 2007) to derive a structural connectome from a DTI scan, then assigned endpoint voxels to regions in the 360-region Atlas to produce one 360×360 SC matrix per individual. We used the same

pipeline as in (Kristanto et al., 2023) to process T1- and T2-weighted structural magnetic resonance imaging scans to compute four features: thickness, myelination, curvature, and sulcus depth. For each feature, we computed the mean over all voxels within a region.

2.2. Ising model simulation and fitting

The generalized Ising model maintained the same formulation used in previous fMRI studies that used them as simplified neural mass models (Ezaki & Watanabe, 2017; Ruffini et al., 2023; Ponce-Alvarez et al., 2022; Theis et al., 2024; Khanra et al., 2024; Theis et al., 2025): Each node i corresponded to a brain region and had an external field parameter, h_i, which influenced its probability of being in the down-regulated state ($\sigma_i = -1$) or up-regulated state ($\sigma_i = +1$). Each pair of regions (i, j) had a coupling parameter, J_{ij}, which influenced the probability of coactivation of two regions and was symmetric ($J_{ij} = J_{ji}$) with no self-coupling ($J_{ii} = 0$). The entropy of a state $\vec{\sigma} = (\sigma_1, \dots, \sigma_N)$ of an N-node system, $s(\vec{\sigma})$, and the partition function L taken over all 2^N possible states determined the probability $P(\vec{\sigma})$:

$$s(\vec{\sigma}) = \sum_{i=1}^{N} \left(h_i \sigma_i + \sum_{j}^{N} J_{ij} \sigma_i \sigma_j \right), L = \sum_{l=1}^{2^N} e^{s(\vec{\sigma}_{(l)})}, P(\vec{\sigma}) = \frac{e^{s(\vec{\sigma})}}{L} \ (1)$$

The most direct method of fitting an Ising model to a sampling of states derived from data is likelihood maximization, but this requires $O(2^N)$ time per optimization step, as one must calculate the probabilities of all possible states (Ezaki & Watanabe, 2017). The most common alternative is pseudolikelihood maximization, which takes $O(TN^2)$ time per optimization step, where T is the number of observed states in the data, but this method only approximates true likelihood maximization for large T (Ezaki & Watanabe, 2017). Consequently, the studies that use it aggregate data from many individuals, as seen in (Ezaki & Watanabe, 2017; Ruffini et al., 2023; Ponce-Alvarez et al., 2022). For examples of results of applying the pseudo-likelihood maximization method from (Ezaki & Watanabe, 2017) to binarized group-level HCP data, see Supplementary Figure S2.

We instead fitted models using the version of Boltzmann learning described in (Roudi et al., 2009). In this method, parameter updates are proportional to the discrepancies between the means $\langle \sigma_i \rangle_{\text{data}}$ and uncentered covariances $\langle \sigma_i \sigma_j \rangle_{\text{data}}$ from the target data and those of the model, $\langle \sigma_i \rangle_{\text{model}}$ and $\langle \sigma_i \sigma_j \rangle_{\text{model}}$ as sampled with Metropolis Monte Carlo simulation (Roudi et al., 2009). The updates to the model parameters are the differences between the model and data first and second order mean fields scaled by an arbitrary learning rate $\eta \epsilon [0,1]$ (Roudi et al., 2009):

$$\delta h_i = \eta(\langle \sigma_i \rangle_{\text{data}} - \langle \sigma_i \rangle_{\text{model}}), \delta J_{ij} = \eta\left(\langle \sigma_i \sigma_j \rangle_{\text{data}} - \langle \sigma_i \sigma_j \rangle_{\text{model}} \right) \ (2)$$

Metropolis simulation was the most computationally costly part of this fitting method, so we used two strategies to make our implementation more efficient. First, we used a sequential update scheme, meaning that we iterated over all the nodes in a fixed order, giving each node state a chance to update, before recording the system state as a new sample. As shown in (Potter & Swendsen, 2013), sequential update maintains total balance while also being the computationally simplest update scheme. In their tests, this allowed it to achieve the fastest convergence to the true probability distribution of states in terms of real-world computation time (Potter & Swendsen, 2013). Second, we used the PyTorch (v1.13.1) framework to run multiple simulations in parallel on a graphics processing unit (GPU). While the Metropolis algorithm requires that we update node states serially, it is straight-forward to run multiple independent simulations of equal-sized systems in parallel using vector arithmetic operations on multi-dimensional arrays.

In addition to simple improvements in computational efficiency, we improved the performance of the Boltzmann learning method by using the mean field values from the data as parameters in an initial guess model and using this initial guess to optimize the simulation temperature. In early tests, we started from a randomly generated model but found that fitting progressed slowly if at all (For an examination of the effect of simulation temperature on the behavior of the random Ising model, see Supplementary Figure S3). By contrast, our initial guess model often started with a correlation between model and binarized data FC close to 0.6, which made it possible for the subsequent fitting to converge more quickly.

While (Roudi et al., 2009) provides several methods for finding approximate fits to data that one can further improve through Boltzmann learning, we found that the even simpler approach of using the means and uncentered covariances derived from the data provided adequate performance (See Figure 3a):

$$h_i = \langle \sigma_i \rangle_{\text{data}}, J_{ij} = \langle \sigma_i \sigma_j \rangle_{\text{data}} \ (4)$$

Intuitively, such a model should behave at least qualitatively similarly to the fitted model. Larger positive values of h_i should increase the probability that region i will remain in the excited state, increasing $\langle \sigma_i \rangle_{\text{model}}$, while negative values of h_i should decrease the probability, decreasing $\langle \sigma_i \rangle_{\text{model}}$. Similarly, larger positive values of J_{ij} should increase the uncentered covariance between regions i and j, $\langle \sigma_i \sigma_j \rangle_{\text{model}}$, while larger negative values should make them anti-correlated, making $\langle \sigma_i \sigma_j \rangle_{\text{model}}$ negative. While this approach worked for both group and individual data, we found that using the group model as the initial guess for the individual models lead to faster convergence and better worst-case fits (Supplementary Figure S4), so we used this group-to-individual approach for the individual models presented in subsequent results.

However, in practice, we found that the FC from the Metropolis simulation would only correlate with that of the data when we simulated near the critical inverse temperature (β). Researchers have long known that simulating at the critical temperature, also known as the Curie temperature, leads to optimal sampling of the model states (Bian et al., 2010). For all choices of binarization threshold, we found that most regions in the binarized data were positively correlated, making most initial $J_{ij} > 0$. Consequently, as β increased (temperature decreased) from 0 to some finite freezing point, the model passed from the fully random phase, through a critical phase with frequent phase transitions but with a clear preference for states that align strongly coupled regions, to a frozen phase in one highly aligned configuration.

Our tests showed that the initial guess model only had an FC correlated with that of the data in a relatively narrow window around the critical β value (See Figure 3a). To optimize the simulation temperature, we performed a 1-dimensional grid-based search, iteratively homing in on the temperature that produced the lowest root mean squared error (RMSE) between the centered covariances of the initial guess model and the target data. We used the covariance RMSE because, even though it is harder to interpret when comparing goodness of fit of different fitted models, it gives numerically valid results even for frozen or near-frozen phases. By contrast, FC correlation can be undefined when regions have 0 variance (no flips). We first searched over the range [0, 1] for each initial guess group model, sampling 101 evenly spaced values per iteration.

For the simulate-and-update loop of Boltzmann learning, we used 1200-step simulations, as longer simulations did not improve convergence time or final goodness of fit of the model. However, we did find that longer simulations of 120,000 steps were necessary for simulated FC to converge to a consistent value. To test a model, we computed the Pearson correlation between its FC and the FC of the target binarized data, where we defined the FC as the Pearson correlation between the states of pairs of regions. We ran these longer test simulations once every 1000 parameter updates to identify the point at which fitting ceased to increase the FC correlation.

2.3. Predictability of model parameters from brain structural features

After fitting Ising models to data, we used least squares linear regression models to test for linear relationships between structural features of the brain and Ising model parameters. We considered four kinds of linear models:

1. Group model J_{ij} predicted from the group mean SC (1 correlation per model).
2. Group model h_i predicted from the group means of the region features (thickness, myelination, curvature, and sulcus depth) (5 correlations per model, one for each feature as the sole predictor and one using all four features).

3. Individual model J_{ij} of a given pair of brain regions from individual SC ($360 \cdot (360 - 1)/2 = 64{,}620$ correlations per model, one for each region pair).
4. Individual model h_i of a given region from individual region features ($360 \cdot 5 = 1800$ correlations per model, 4 single-predictor regressions and one 4-predictor regression per region).

To demonstrate the value of using the model parameters instead of comparing structural feature data directly to fMRI data, we compared each linear regression model to a similar model predicting a feature derived directly from the binarized fMRI data, either the mean state of the binarized region time series for h_i or the FC of the binarized time series for J_{ij}. We measured the goodness of each linear model using the correlation between predicted-from-structure and actual Ising model parameter values.

Finally, we looked for spatial trends in the local standard deviation (SD) of model parameters and the strength of the local individual-level prediction correlations. Hypothesizing that a lack of individual variation in brain structure might explain why some regions had poor predictability of individual model h_i from structure, we computed the correlations between the local SD of each model parameter and the local prediction correlation. Similarly, we computed correlations between the local SD of SC and the prediction correlation of individual model J_{ij}. Since several lines of research, including those described in (Fotiadis et al., 2023) and (Paquola et al., 2019), have suggested that regions with greater myelination trade more efficient signal transmission for lower synaptic plasticity or functional variability, we tested whether mean myelination negatively correlated with individual variability (SD) of h_i or positively correlated with predictability of h_i from structure.

To test each prediction correlation for statistical significance, we used permutation tests in which we repeated the linear regression after permuting the feature-parameter pairs. For individual model tests, we used $\alpha = 0.05$ and applied Bonferroni correction using the number of regions or of region pairs as appropriate. For group model correlations and other correlations with group level mean or SD of structural features, we selected $\alpha = 0.001$ without an explicit Bonferroni correction. To avoid biasing the tests by having multiple replicas of the same model, we took the mean of each parameter over replicas. In the Results section, we use $p < 10^{-6}$ to denote that none of the 10^6 permuted correlations were stronger than the original.

3. Results

3.1. Effect of binarization of data on functional connectivity

A crucial initial question is how binarization impacts the FC as estimated from the data itself. Binarization is necessary to create an achievable fitting target, but it may distort the FC. We searched for the interval of thresholds over which binarization preserved the

overall pattern of FC. Figure 2 shows how binarizing the data affects activity in an example time series (a), alongside covariance (b) and FC (c). Comparing (b) and (c) showed that covariance diminished due to the diminishing variances of the individual region time series, while FC, which was normalized with respect to variance, remained in-tact. For these plots, we selected a maximum threshold of 1.6 because this was the largest value at which every brain region of every individual scan time series had a non-0 variance. Plot (d) illustrates this correspondence between the original and binarized example time series FC values of individual region pairs at a binarization threshold of 1, for which the correlation is 0.93. Plot (e), compares FC correlations between original and binarized time series data at binarization thresholds from 0 to 2.5. The median (black line) and range (shaded area) of individuals show that binarization reliably preserved individual FC up to a threshold of 1, whereas the FC correlation for the group aggregate data (red line) remained high up to a threshold of 2. At threshold 1, the worst case single-subject FC correlation was 0.88.

Figure 2. Effect of binarization on FC of fMRI time series data (**a**) First half (432 seconds) of an example fMRI time series from a randomly selected HCP subject. The top image shows the preprocessed, parcellated, and z-scored but un-binarized values. Subsequent images show the same time series binarized at thresholds 0, 1, and 1.6. (**b**) Covariances between pairs of regions in the original and binarized region time series. (**c**) FC (Pearson correlation) between pairs of regions in the original and binarized time series using all time points of all scans from the

individual subject. (**e**) Scatter plot showing relationship between original and binarized FC of all region pairs for this subject. (**d**) Correlation between original and binarized time series FC as a function of binarization threshold. The black line shows the median, and the shaded area the range of individual subject FC correlations. The red line shows the FC correlation of group aggregate data (concatenated z-scored time series of 670 individuals).

3.2. Goodness-of-fit of models fitted to group and individual data

Having determined the range of thresholds over which binarization preserves the FC of the data, we next tested our Ising model fitting workflow with the same data binarized at different thresholds. We also investigated the importance of the h_i parameters to model behavior. We found that, as the threshold increased, the initial guess model became less effective as an approximation of the fitted model, leading to slower and less reliable convergence, but, within the range where fitting succeeded, larger thresholds caused the heterogeneity of h_i play a more significant role in the simulation.

As part of our fitting process, we needed to identify the optimal simulation temperature for each threshold. Due to the computational cost of fitting a full set of individual models, we tested the same 670-subject pooled group data binarized at 26 different thresholds from 0 to 2.5. Figure 3a shows that, prior to fitting, each initial guess group model had an optimal simulation temperature that maximizes FC. However, as the threshold increased, the peak in FC correlation shrank, disappearing near 2.0.

Because Boltzmann learning relies on the stochastic Metropolis algorithm to sample states from the model, we tested the reliability of fitting by independently fitting 101 model instances to each binarized group data target. After 40,000 parameter updates, fitting had converged to 0.95 or more for all models fitted to group data binarized at thresholds less than or equal to 1 but not at higher thresholds, such as 1.6 (Figure 3b). Past the threshold 1, the final FC correlation dropped sharply (Figure 3c). Since binarization still preserved the group level FC in this range (see Figure 2e), FC correlation between the model and binarized data (blue line) matched FC correlation between the model and un-binarized data (orange line), even as both declined. We also tested for convergence of the underlying model parameters and found that replica models remained nearly identical at thresholds up to 0.8 and remained mostly consistent except for some outlier replicas at higher thresholds (Supplementary Figure S5).

Since we intended for the Ising models to serve as a bridge not only between FC and SC but also between local differences in brain structure and local differences in brain dynamics, we needed the heterogeneity of the local parameter, h_i, to significantly affect model behavior. We tested the influence of h_i by comparing the FC correlation of each fitted group model to that of a modified version in which we replaced the heterogeneous

h_i values with their mean. At thresholds close to 0, homogenizing h_i in this way had no effect, but, at thresholds close to 1, removing the heterogeneity clearly worsened the fit between model and data (Figure 3d). We also found that, whereas the structure of the J_{ij} matrix remained similar and retained similar levels of heterogeneity at all thresholds, h_i vectors had similar patterns of heterogeneity across a wide range of positive thresholds but not at 0 (Supplementary Figure S6).

Having established the range of thresholds within which we could fit the Ising model to pooled group data, we next fitted models to individual data and tested whether the heterogeneity of h_i played any role in reproducing individual differences in FC. We selected two thresholds near the ends of the range where group model fitting had succeeded, 0 and 1. To test the reliability of individual model fitting, for each threshold, for each of the 837 subjects, we made 5 replica models. We found that models for threshold 1 took longer to converge, with FC correlation only ceasing to improve substantially after 40,000 parameter updates, as opposed to only 10,000 for threshold 0. Furthermore, whereas fitting at threshold 0 achieved FC correlations of 0.90 or more, fitting at threshold 1 yielded a small fraction of failed fitting attempts with FC correlations as low as 0.34. These failed fitting attempts did not appear to result from issues with any one subject's data, as every subject had at least 3 replicas with FC correlations above 0.90. They did, however, appear to be genuine failures of model fitting instead of failures of the test simulation to converge to a stable FC, as the same replica models had low FC correlations across multiple tests. We also checked that the FC correlations between individual models and individual data were significantly higher than the FC correlations between group models and individual data (1-tailed Wilcoxon signed rank test $p < 0.001$ for both thresholds), showing that the individual models accurately captured individual differences in FC (Figure 3e). We also found that the FC correlations of the 670 subjects pooled for the group model, which served as the initial guess for the individual model fitting, were not significantly greater than those of the remaining 167 subjects (1-tailed Mann-Whitney U-test p=0.97), showing that the group model could serve as a good initial guess even for new individual subjects (Supplementary Figure S4). Since the violin plots of FC correlations for in-group and out-group subjects were visually identical to the plot for all subjects, we have omitted them here.

Having established that region-to-region heterogeneity of h_i affects model behavior at high binarization thresholds but not at low ones and that individual models capture individual differences in FC, we next tested the importance of subject-to-subject differences in h_i to the ability of the individual models to capture individual differences. To do this, we compared the FC correlations of the fitted individual models to FC correlations of modified versions of the individual models in which we replaced each h_i with the group mean value for region i. We found that removing individual differences in h_i did not significantly decrease FC correlation for threshold 0 (1-tailed Wilcoxon signed

rank test p=0.17) but did for threshold 1 (1-tailed Wilcoxon signed rank test p<0.001) (Figure 3f). We also examined the underlying model parameters and found that, whereas J_{ij} of a given pair (i, j) had nearly the same mean and SD over individuals at either threshold, mean h_i values did not correlate between thresholds 0 and 1, while the SD over individuals only weakly correlated while typically being twice as large for threshold 1 (Supplementary Figure S7), showing that the higher threshold had induced a greater degree and different pattern of heterogeneity among nodes.

Figure 3: Goodness-of-fit measured by correlation between model and binarized data FC. (**a**) Effect of inverse temperature (beta) on correlation between the FC of binarized group-level data and the FC of the data-derived initial guess Ising model. Each curve shows results for a different binarization threshold, 0, 1, or 1.6. (**b**) Progression of group model fitting measured by correlation between binarized data FC and model FC. Lines represent medians, and shaded areas represent ranges over 101 independent fittings to the same data. (**c**) Correlations between FC of

fitted group models with either binarized group data (blue) or with continuous-valued group data (orange) at a range of binarization thresholds. The lines represent medians, and shaded areas represent ranges over 101 independently fitted replica models. (**d**) Correlation between the FC of the group model and binarized group data for the complete fitted model, including heterogeneous h_i, (blue) and for a modified model using the mean value of h_i at all regions (orange). (**e**) Comparison of the distributions of group model FC versus individual data FC correlations and distributions of individual model FC versus individual data FC correlations. The number at the end of the label indicates the binarization threshold, 0 or 1. Ticks show the quartiles. Curly braces show pairs of distributions that are significantly different at the 0.001 level (no Bonferroni correction). (**f**) Comparison of distributions of FC correlations for individual model versus individual data for full individual models, including fitted h_i, and for modified individual models with fitted individual J_{ij} but group mean h_i at each region. The number at the end of the label indicates the binarization threshold, 0 or 1. Ticks show quartiles. The curly brace and asterisk indicate that the difference for threshold=1 is significant (p < 0.001).

3.3. Predictability of model parameters from structural features

Because the Ising model fitting process disentangles the local excitability, represented by h_i, from the connectivity between regions, represented by J_{ij}, we expected that Ising model parameters could represent the underlying characteristics of brain regions better than could statistical properties calculated directly from the data, such as local mean state and FC. Works such as (Ruffini et al., 2023) have already reported correlations between J_{ij} and SC, but we also hypothesized that h_i would correlate with local structural features that influence excitability. We compared parameters and structural features at two levels: brain-wide correspondences between group model parameters and group mean brain structure, and local correspondences between individual differences in model parameters and individual differences in brain structure.

First, for each of the Ising models that we fitted to group data binarized at a different threshold, we fitted a least squares linear regression model $\hat{J}(SC_{ij})$ to group model J_{ij} and group mean SC_{ij} across all region pairs (i, j). We then computed the correlation between J_{ij} and $\hat{J}(SC_{ij})$ (Figure 4a, blue line). For comparison, we computed the correlation between FC_{ij} of the binarized group time series data and the linear regression model $\widehat{FC}(SC_{ij})$ (Figure 4a, orange line). While J_{ij} correlation dropped and FC_{ij} correlation rose as the threshold increased, both remained well above their permutation test critical values (Figure 4a, broken line), and J_{ij} correlation remained stronger at all thresholds, showing that J_{ij} can serve as a better bridge between structural and functional brain networks.

As a further test of the robustness of these correlations, we cross-validated them using 10,000 random splits of region pairs into equal-sized training and testing sets. At all

thresholds, training and testing correlations were largely identical, both forming a narrow band centered on the correlation for the full set of region pairs (Supplementary Figures S8 a and b). We also used the linear models that we fitted to group data to predict individual J_{ij} and FC_{ij} from SC_{ij} and found that, while the correlations (1 per subject) were lower, correlations were still consistently higher for J_{ij} than for FC_{ij} and were similar regardless of whether the subjects were included in the group model (Supplementary Figure S9a) or not (Supplementary Figure S9b).

Having established that increasing the binarization threshold increased the heterogeneity of h_i and the effect of that heterogeneity on the goodness of fit of the Ising models (Figure 3, plots d and f), we next hypothesized that higher thresholds would also be necessary to observe correlations between h_i and brain structural features. To test this, we compared group model h_i at multiple thresholds to group mean structural MRI features. To integrate information from all four features, for each threshold, we fitted a least squares multiple linear regression model $\hat{h}(t_i, m_i, c_i, d_i)$ to h_i, thickness (t_i), myelination (m_i), curvature (c_i), and sulcus depth (d_i). As expected, we found that the correlation between h_i and $\hat{h}(t_i, m_i, c_i, d_i)$ trended upward with increasing threshold (Figure 4b, blue line), exceeding the permutation test critical value ($\alpha = 0.001$, Figure 4a, broken line) when the threshold reached 0.3.

To test the relative importance of different features, we compared these results to direct correlations with individual features and with single-feature linear regression models (equivalent to the absolute value of the direct correlation) and found that correlation with $\hat{h}(t_i, m_i, c_i, d_i)$ closely tracked and was only marginally better than correlation with the myelination alone (Figure 4b, green line). The direct correlation between h_i and myelination was positive at all thresholds, indicating that more myelinated regions differentially activated more readily. All four direct h_i-feature correlations became stronger with increasing threshold, with curvature being the only feature not to achieve significance before threshold 1.0 (Supplementary Figure 10a).

For comparison, we also computed the correlation between the mean state of the binarized group data, $\langle \sigma_i \rangle$, and the linear model predicting it from structural features, $\widehat{\langle \sigma \rangle}(t_i, m_i, c_i, d_i)$ (Figure 4b, orange line). We found that $\langle \sigma_i \rangle$ correlation changed non-monotonically with threshold, starting out higher than h_i correlation, dipping to a minimum near 1.1, then rising again but remaining below h_i correlation. As with h_i correlation, the four-feature $\langle \sigma_i \rangle$ model correlation closely tracked the absolute value of the single-feature correlation with myelin, but, in this case, that correlation flipped from negative to positive between thresholds 1.1 and 1.2 (Supplementary Figure S10b). Overall, the maximum correlations for h_i (0.55) and $\langle \sigma_i \rangle$ (0.45) were close to each other, suggesting that both vary along whole-brain-scale gradients in structural features.

We cross-validated the 4-feature group-level correlations using random, even splits of brain regions and found that, while testing set correlations were consistently slightly lower than training set correlations, they remained close and reproduced the same trends with respect to binarization threshold (Supplementary Figures S8 c and d). We also applied the group-level 4-feature linear models to individual subject data and found that, for each type of correlation, the individual correlations (1 per subject) covered a much broader range with the group-level correlations at the upper end of it (Supplementary Figures S9 c and d). However, for both training and testing subjects, h_i with threshold 1 generalized better to individuals compared to h_i at threshold 0 or $\langle \sigma_i \rangle$ at either threshold.

We next tested whether individual differences in J_{ij} corresponded to individual differences in SC. For each region pair (i, j), we fitted a separate linear model $\hat{J}_{ij}\left(SC_{ij}^{(s)}\right)$ to individual model $J_{ij}^{(s)}$ parameters and $SC_{ij}^{(s)}$ data values of all subjects s. For either binarization threshold (0 or 1), we compared the distribution of correlations between Ising model $J_{ij}^{(s)}$ and $\hat{J}_{ij}\left(SC_{ij}^{(s)}\right)$. We also compared these to correlations between individual binarized data FC, $FC_{ij}^{(s)}$, and FC predicted from SC, $\widehat{FC}_{ij}\left(SC_{ij}^{(s)}\right)$. Figure 4c shows these four distributions, each consisting of 64,620 correlations.

In all four cases, most of these local individual difference correlations were much weaker than the global correlation between group model and group mean data (Figure 4a), and fewer than 1% were statistically significant (1-tailed permutation test with $\alpha = \frac{0.05}{64,620} \approx 7.7 \cdot 10^{-7}$). $J_{ij}^{(s)}$ at threshold 0 had the most significant correlations (602), followed by $J_{ij}^{(s)}$ at threshold 1 (495), then $FC_{ij}^{(s)}$ at threshold 0 (360), and finally $FC_{ij}^{(s)}$ at threshold 1 (310). At either threshold, the significant $J_{ij}^{(s)}$ correlations tended to be larger than the significant $FC_{ij}^{(s)}$ correlations (1-tailed Mann-Whitney U-test, both $p < 10^{-7}$). This suggests that, while the individual difference correlations were mostly weak, $J_{ij}^{(s)}$ prediction correlation can serve as a more sensitive measure with which to find those connections between brain regions that do have strong structure-function coupling. We confirmed the robustness of these findings by cross-validating with 100 random 670-167 (nearly 80-20) splits of the subjects into training and testing sets. While correlations that were weak on average tended to vary widely across different splits, the stronger correlations were more consistent, with average training and testing correlations matching closely (Supplementary Figures S11 a-d).

Similarly, we tested whether local differences in structural features predicted local differences in h_i. For each region i, we fitted a linear regression model $\hat{h}_i\left(t_i^{(s)}, m_i^{(s)}, c_i^{(s)}, d_i^{(s)}\right)$ to individual Ising model $h_i^{(s)}$ and structural feature values and

computed the correlation between it and the Ising model $h_i^{(s)}$ over all subjects s. We compared results for thresholds 0 and 1 and compared these to correlations between $\langle \sigma_i \rangle^{(s)}$ from the binarized data and linear regression models $\widehat{\langle \sigma_i \rangle}\left(t_i^{(s)}, m_i^{(s)}, c_i^{(s)}, d_i^{(s)}\right)$. Figure 4d shows these four distributions, each consisting of 360 correlations.

Again, in all four cases, only a minority of regions had significant prediction correlations, but, unlike in the comparison of individual-level $J_{ij}^{(s)}$ and $FC_{ij}^{(s)}$ correlations, where the distributions were visually similar, necessitating selection of significant correlations and statistical testing to determine which ones tended to be larger, $h_i^{(s)}$ prediction correlations at threshold 1 were clearly higher than $h_i^{(s)}$ prediction correlations at threshold 0 or $\langle \sigma_i \rangle^{(s)}$ correlations at either threshold. Furthermore, this also applied to direct correlations with individual structural features (Supplementary Figures S8 c-f). Unlike with group-level correlations, no one structural feature was consistently the strongest predictor of $h_i^{(s)}$ (Supplementary Figure S8h). For $\alpha = \frac{0.05}{360} \approx 1.4 \cdot 10^{-4}$, 96 threshold 1 $h_i^{(s)}$ correlations, 12 threshold 0 $h_i^{(s)}$ correlations, 5 threshold 1 $\langle \sigma_i \rangle^{(s)}$ correlations, and 9 threshold 0 $\langle \sigma_i \rangle^{(s)}$ correlations were significant. This result shows that $h_i^{(s)}$ has a stronger relationship to local brain structure than does $\langle \sigma_i \rangle^{(s)}$, but only at thresholds above 0. Cross-validation showed that, while weaker correlations were consistently lower over the testing set subjects than the training set subjects, this difference narrowed as the average correlations grew stronger (Supplementary Figures S11 e-h).

We also searched for specific factors that could influence which connections had stronger predictability of individual J_{ij} from SC or which regions had stronger predictability of individual h_i from structural features. For these analyses, we focused on individual Ising models at threshold 1. We first hypothesized that weak correlations could be due to a lack of significant individual variation in structure. To test this, we calculated the Pearson correlations between the individual-level prediction correlation and the SD of each feature over individuals. We found that J_{ij} prediction correlation moderately but significantly correlated with the SD of SC (37%, 1-tailed permutation test $p < 10^{-6}$, Figure 4e) and with SD of J_{ij} (25%, $p < 10^{-6}$), suggesting that connections with little variability of SC and J_{ij} account for at least some of the low prediction correlations.

Of the four structural MRI features, the $h_i^{(s)}$-$\hat{h}_i\left(t_i^{(s)}, m_i^{(s)}, c_i^{(s)}, d_i^{(s)}\right)$ correlation correlated most strongly with SD of curvature (29%, $p < 10^{-6}$, Figure 4f). The correlations with the SD of thickness and SD of myelination were not significant. While the correlation with SD of sulcus depth was significant (also 29%, $p < 10^{-6}$), the SD of curvature and SD of sulcus depth correlated strongly with each other (71%), suggesting that they mostly provided redundant information. Cortical surface maps of $h_i^{(s)}$ prediction

correlation (Figure 4h) and of the SD of curvature (Figure 4i) show that, while the regions with the strongest structure-function couplings and those with the most variable curvature were not always the same, they tended to occur close to each other.

Comparing $h_i^{(s)}$ prediction correlations of regions and their descriptions in (Glasser et al., 2016) showed that higher visual areas had a particularly high concentration of high correlations. Of the top 36 regions (10%), 6 belonged to the ventral stream visual cortex, 6 to the MT+ complex and associated occipital and temporal visual areas. Another 3 were part of the nearby temporal-parietal-occipital junction, which bridges higher visual and auditory regions, and 2 were in the adjacent lateral temporal cortex. By contrast, the primary and early visual areas were all in the bottom half of the ranking. Most of the remaining top 10% were in the frontal lobe, including the premotor areas (5 regions), anterior cingulate and medial prefrontal cortex (3), inferior frontal cortex (3), orbital and polar frontal cortex (2), dorsolateral prefrontal cortex (1).

As a further test of the biological relevance of individual differences in h_i, we tested whether we could use them to observe that structure-function correlation decreased and synaptic plasticity increased along the gradient from highly myelinated unimodal sensory and motor areas to minimally myelinated high-level association areas as observed by (Paquola, 2019). We first tested the correlation between 4-feature $h_i^{(s)}$ correlation and mean myelination. While the trend was positive, it was not significant ($0.01, p = .85$). However, the local SD of $h_i^{(s)}$ negatively correlated with mean myelination (-52%, $p < 10^{-6}$, Figure 4g). Specifically, the cortical maps of both mean myelination (Figure 4j) and SD of h_i (Figure 4k) show opposite trends along the axis from the highly myelinated sensory, motor, and occipital cortices to the less myelinated frontal and temporal areas.

Figure 4: Correlations between fitted Ising model parameters and brain structural features. (**a**) Effect of binarization threshold on the correlation between group model J_{ij} and mean SC (blue) and the correlation between binarized group data FC and mean SC (orange). The broken line represents the largest critical value corresponding to $\alpha = 0.001$ from any of the permutation tests of significance, though all were close to 0.01. (**b**) Effect of threshold on correlations between

group Ising model h_i and h_i as predicted from group mean structural MRI features (thickness, myelination, curvature, and sulcus depth) with a multiple least squares regression model (blue), h_i and h_i as predicted from mean myelination alone (green), and binarized mean state of the group-level data versus mean state predicted from the four structural features (orange). The broken line represents the largest critical value for $\alpha = 0.001$ in any of the permutation tests, though all were approximately 0.17. (**c**) Distribution of correlations between individual variations in individual Ising model J_{ij} and J_{ij} as predicted from SC with a least squares linear regression model (one linear model and correlation for each pair of brain regions) or between FC of the binarized data and FC predicted from SC. (**d**) Distribution of correlations between local variations in individual Ising model h_i and h_i as predicted from the four structural features with a least squares linear regression model (one linear model and correlation for each brain region) (**e**) Scatter plot of correlation between individual Ising model J_{ij} (threshold=1) and J_{ij} predicted from SC versus SD of SC (one point per pair of brain regions). The orange line shows the least squares regression trend. (**f**) Scatter plot of correlation between individual Ising model h_i (threshold=1) and h_i predicted from the four structural features versus SD of curvature (one point per brain region). The orange line shows the least squares regression trend. (**g**) Scatter plot of SD of individual Ising model h_i (threshold=1) versus mean myelination (one point per brain region). The orange line shows the least squares regression trend. (**h**) Cortical surface map of correlation between individual Ising model h_i (threshold=1) and h_i predicted from the four structural features. (**i**) Cortical surface map of SD of curvature. (**j**) Cortical surface map of correlation between individual Ising model h_i (threshold=1) and h_i predicted from the four structural features. (**j**) Cortical surface map of SD of individual Ising model h_i (threshold=1). (**k**) Cortical surface map of mean myelination.

4. Discussion

Overall, our work presents a more comprehensive approach to using Ising models as simplified, data-driven neural mass models to study the relationship between brain structure and functional characteristics. We started by introducing methodological improvements to enable Boltzmann learning to scale to larger, more fine-grained representations of the brain. We then studied how data binarization threshold influenced the role of the heterogeneity of h_i in allowing the model to reproduce the FC of the data. Finally, we studied how binarization threshold affected relationships between J_{ij} and SC and between h_i and brain region structural features.

4.1. An improved approach to fitting to fMRI data

Fitting the group-level model by starting with a data-derived initial guess model, optimizing the simulation temperature prior to fitting the parameters, and using the group model as the initial guess for individual-level models lead to faster convergence and greater scalability (Supplementary Figure S7). Whereas prior studies that fitted Ising models to fMRI data fitted them only to group-level data using pseudolikelihood

maximization (Ezaki & Watanabe, 2017; Fortel et al., 2022; Ponce-Alvarez et al. 2022; Ruffini et al., 2023) or only fitted small models of 10 or fewer nodes to individual data using likelihood maximization (Watanabe et al., 2013; Theis et al., 2024; Theis et al., 2025) or Bayesian approximations (Jeong et al., 2021; Kang et al., 2021; Khanra et al., 2024), our approach can fit fully connected Ising models to a 360-region parcellation of individual-level fMRI data, allowing more precise identification of key regions and connections of importance to individual differences in patterns of whole-brain integration and segregation.

Furthermore, our study is the first to address how choice of binarization threshold affects the role of node heterogeneity in the Ising model, which is important to comparisons between the Ising model and the real brain, in which brain regions have distinct intrinsic dynamical properties. Of the prior studies that used fMRI data as a fitting target, only (Theis et al., 2025), (Fortel et al., 2022), and (Watanabe et al., 2013) compared Ising models fitted to data binarized at multiple thresholds. All three reported that binarization at the mean produced the closest fit between model and data. While this also held true for our method, it achieved good correlations among model simulation FC, binarized data FC, and continuous-valued data FC over a range of thresholds from 0 to 1 SD above the mean BOLD signal.

Whereas the above-cited works did not systematically examine the effects of binarization threshold on the underlying parameters of the fitted model, we showed that it influenced the extent to which the model was merely a network of interchangeable nodes or a network of nodes with distinct functional characteristics. At a threshold of 0, removing region-to-region heterogeneity of h_i did not affect the FC correlation of the group model, and removing subject-to-subject heterogeneity of h_i did not affect the FC correlation of the individual models. This implies that the nodes were effectively homogeneous and that J_{ij} contained all information needed to reproduce the FC of the data. By contrast, using a threshold of 1 maintained good FC correlation and induced both the group and individual models to use heterogeneous h_i values that played significant roles in reproducing the FC of the original data. This better reflects the real human brain, in which distinct regions have different morphology and in which individual differences in local region structural features correlate substantially with excitability (Glasser et al., 2014; Fotiadis et al., 2023; Fotiadis et al., 2024), which in turn influences larger-scale brain dynamics (Abeywardena, 2024).

4.2. An intermediary between structural and functional connectivity

To test the biological relevance of the fitted model parameters, we computed correlations between J_{ij} and SC and between h_i and four local region features: thickness, myelination, curvature, and sulcus depth. We found that, for any FC-preserving binarization threshold, the least squares regression model prediction from group mean SC correlated better with

group model J_{ij} than with group binarized data FC, showing that the Ising model can serve as an intermediary between the structural network and the functional network (Figure 4c). As suggested in (Ruffini et al., 2023), J_{ij} could serve as a signed approximation of the structural connectome in future modeling studies with more sophisticated models such as neural mass models, where recovering the signed structural connectivity from functional data is more difficult.

While earlier studies have found similarly robust relationships between population-level SC and J_{ij} (Watanabe et al., 2013; Fortel et al., 2022; Ruffini et al., 2023), our approach allowed us to study the relationship between individual differences in local SC and J_{ij} in a way that other studies have not. While individual-level correlations were much more variable and generally much weaker, the number and strength of significant SC-J_{ij} correlations were significantly greater than for FC-SC correlations (Figure 4c and Supplementary Figure S10g). Furthermore, the largest SC-J_{ij} correlations occurred on connections with greater variability of SC (Figure 4e), suggesting that the individual Ising models do succeed at capturing the most significant differences in underlying connectivity. These results indicate that, in future studies, the J_{ij} can serve as an intermediate representation through which to link individual differences in fMRI-derived FC and DTI-derived SC more effectively than direct correlations between the two.

4.3. A mapping from structural features to local excitability

Whereas researchers have previously noted the correspondence between SC and J_{ij} (Watanabe et al., 2013; Ponce-Alvarez, 2022; Ruffini et al., 2023), none have studied the relationship between h_i and brain structure. Whereas mean activation of a region depends on both its intrinsic excitability and its connectivity with other regions, the Ising model fitting process disentangles these aspects of local behavior, capturing connectivity in J_{ij} and intrinsic excitability in h_i. This suggests that h_i can serve as a representation of local excitability and that correlates with the underlying physical features of the brain that give rise to differences in excitation-inhibition balance.

Combining multiple such features by using a multiple linear regression model produced better predictions of h_i, showing that multiple aspects of brain structure influence activity. As with the SC- J_{ij} correlations, we considered both region-to-region heterogeneity based on pooled population-level data and local individual differences. While group-level h_i prediction correlations were robust, they were only slightly stronger than the best group-level binarized mean state prediction correlations (Figure 4b). By contrast, h_i clearly achieved stronger correlations with individual differences in structure (Figure 4d). This shows that h_i is especially valuable for representing individual rather than group differences in excitability and can serve as a more sensitive measure of structure-function coupling.

To develop a better sense of how specific structural features influenced these differences, we searched for region-to-region trends in structural features that correlated with the individual-level variability or predictability of h_i. At the whole-brain level, myelination was the dominant predictor of group-level h_i (Supplementary Figure S10a), which increased along the axis from less myelinated multimodal information integration regions to unimodal sensory and motor regions. This is consistent with the hypothesis that myelination offers increased efficiency of signal transmission, represented in the Ising model by easier induction of differential activation due to higher h_i. However, along this same axis, we saw a decrease in individual variation in h_i (Figure 4g), suggesting this increased efficiency comes at the cost of reduced synaptic plasticity. While prior studies such as (Paquola, 2019) and (Fotiadis et al., 2023) have argued for this tradeoff based on mechanistic hypotheses and correlations between group-level functional and structural gradients, our results show that h_i can provide a measure of both spatial group-level functional gradients and individual-level differences derived from fMRI data by disentangling excitability (measured by h_i) from connectivity (measured by J_{ij}).

However, variability of curvature, not myelination, had the strongest relationship with the predictability of h_i (Figure 4f). Furthermore, combining information from multiple features offered a substantial advantage when predicting individual differences in h_i (Compare Figure 4d to Figures S10 c-f). This suggests that, whereas myelination constrains the capacity for individual variation, the actual differences in brain dynamics arise through changes in gray matter cortical architecture, particularly through differences in the amount of cortical folding. This is consistent with an earlier study, which found that sulcus depth correlated with individual variation in the resting state functional connectivity between a region and surrounding areas (Mueller et al., 2013). That the highest h_i prediction correlations concentrated in the frontal lobe and in higher visual and multimodal areas also suggests that regions with higher synaptic plasticity have more room for changes in architecture that can influence excitability and thus h_i.

Conclusion

The Ising model fitting workflow and parameter analyses presented here represent a substantial step forward in the use of Ising models as simplified, data-driven neural mass models. By fitting models to whole-brain data instead of pre-selecting a set of regions to include, researchers can discover important subnetworks that may be specific to the individual. Furthermore, selecting a threshold that induces node-level heterogeneity will produce models that reflect both the interconnectivity and distinctness of brain regions. This approach will help to bridge the gap between connectomics and translational neuroscience. While connectomics draws on increasingly complex network-level views of the brain (Mišić & Sporns, 2016), clinical research continues to focus on single regions

when designing interventions such as transcranial direct current stimulation (Kang et al., 2024) and transcranial magnetic stimulation (Cash & Zalesky, 2024). Future work will focus on both the discovery of nonlinear and non-local associations between structural features and Ising model parameters and the identification of networks of regions and connections that are important to individual differences in cognitive abilities and mental health.

ACKNOWLEDGMENTS

This work was partially supported by STI 2030-Major Projects (No. 2022ZD0208500), the Hong Kong Research Grant Council Senior Research Fellow Scheme (SRFS2324-2S05) and General Competitive Fund (GRF12202124). We performed most computations for this study on the HKBU DAAI Computing Cluster, which was supported by HKBU Strategic Development Fund.

Author contributions: Conceptualization: CZ; Methodology: AC, SC, QT; Investigation: AC; Visualization: AC, SC, QT; Writing—original draft: AC; Writing—review \& editing: all authors.

Competing interests: Authors declare that they have no competing interests.

Data and materials availability: The MRI data are available at http://www.humanconnectome.org/study/hcp-young-adult. The code used for this study is available on request.

References

Abeyasinghe, P. M., De Paula, D. R., Khajehabdollahi, S., Valluri, S. R., Owen, A. M., & Soddu, A. (2018). Role of dimensionality in predicting the spontaneous behavior of the brain using the classical Ising model and the Ising model implemented on a structural connectome. Brain connectivity, 8(7), 444-455. https://doi.org/10.1089/brain.2017.0516

Abeyasinghe, P. M., Aiello, M., Nichols, E. S., Cavaliere, C., Fiorenza, S., Masotta, O., ... & Soddu, A. (2020). Consciousness and the dimensionality of DOC patients via the generalized ising model. *Journal of clinical medicine*, 9(5), 1342. https://doi.org/10.3390/jcm9051342

Abeywardena, K. G. (2024). *The impact of intrinsic excitability heterogeneity on the macro-scale brain dynamics: Computational analysis on stability, resilience, and synchrony* (Master's thesis, University of Toronto (Canada)). http://hdl.handle.net/1807/140525

Ashourvan, A., Shah, P., Pines, A., Gu, S., Lynn, C. W., Bassett, D. S., ... & Litt, B. (2021). Pairwise maximum entropy model explains the role of white matter structure in shaping emergent co-activation states. *Communications biology*, 4(1), 210. https://doi.org/10.1038/s42003-021-01700-6

Bansal, K., Nakuci, J., & Muldoon, S. F. (2018). Personalized brain network models for assessing structure–function relationships. *Current Opinion in Neurobiology*, 52, 42-47. https://doi.org/10.1016/j.conb.2018.04.014

Behrens, T. E., Berg, H. J., Jbabdi, S., Rushworth, M. F., & Woolrich, M. W. (2007). Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *Neuroimage*, 34(1), 144-155. https://doi.org/10.1016/j.neuroimage.2006.09.018

Bian, Z., Chudak, F., Macready, W. G., & Rose, G. (2010). The Ising model: teaching an old problem new tricks. *D-wave systems*, 2, 1-32.

Bijsterbosch, J., Harrison, S. J., Jbabdi, S., Woolrich, M., Beckmann, C., Smith, S., & Duff, E. P. (2020). Challenges and future directions for representations of functional brain organization. *Nature neuroscience*, 23(12), 1484-1495. https://doi.org/10.1038/s41593-020-00726-z

Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nature neuroscience*, 20(3), 340-352. https://doi.org/10.1038/nn.4497

Cash, R. F., & Zalesky, A. (2024). Personalized and circuit-based transcranial magnetic stimulation: evidence, controversies, and opportunities. *Biological psychiatry*, 95(6), 510-522. https://doi.org/10.1016/j.biopsych.2023.11.013

Das, T. K., Abeyasinghe, P. M., Crone, J. S., Sosnowski, A., Laureys, S., Owen, A. M., & Soddu, A. (2014). Highlighting the structure-function relationship of the brain with the Ising model and graph theory. *BioMed Research International*, 2014(1), 237898. https://doi.org/10.1155/2014/237898

Deco, G., Kringelbach, M. L., Arnatkeviciute, A., Oldham, S., Sabaroedin, K., Rogasch, N. C., ... & Fornito, A. (2021). Dynamical consequences of regional heterogeneity in the brain's transcriptional landscape. *Science Advances*, 7(29), eabf4752. https://doi.org/10.1126/sciadv.abf4752

Demirtaş, M., Burt, J. B., Helmer, M., Ji, J. L., Adkinson, B. D., Glasser, M. F., ... & Murray, J. D. (2019). Hierarchical heterogeneity across human cortex shapes large-scale neural dynamics. *Neuron*, 101(6), 1181-1194. https://doi.org/10.1016/j.neuron.2019.01.017

Domhof, J. W., Eickhoff, S. B., & Popovych, O. V. (2022). Reliability and subject specificity of personalized whole-brain dynamical models. *NeuroImage*, *257*, 119321. https://doi.org/10.1016/j.neuroimage.2022.119321

Elam, J. S., Glasser, M. F., Harms, M. P., Sotiropoulos, S. N., Andersson, J. L., Burgess, G. C., ... & Van Essen, D. C. (2021). The human connectome project: a retrospective. *NeuroImage*, 244, 118543. https://doi.org/10.1016/j.neuroimage.2021.118543

Ezaki, T., Watanabe, T., Ohzeki, M., & Masuda, N. (2017). Energy landscape analysis of neuroimaging data. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2096), 20160287. https://doi.org/10.1098/rsta.2016.0287

Fortel, I., Butler, M., Korthauer, L. E., Zhan, L., Ajilore, O., Sidiropoulos, A., ... & Leow, A. (2022). Inferring excitation-inhibition dynamics using a maximum entropy model unifying brain structure and function. *Network Neuroscience*, 6(2), 420-444. https://doi.org/10.1162/netn_a_00220

Fotiadis, P., Cieslak, M., He, X., Caciagli, L., Ouellet, M., Satterthwaite, T. D., ... & Bassett, D. S. (2023). Myelination and excitation-inhibition balance synergistically shape structure-function coupling across the human cortex. *Nature Communications*, *14*(1), 6115. https://doi.org/10.1038/s41467-023-41686-9

Fotiadis, P., Parkes, L., Davis, K. A., Satterthwaite, T. D., Shinohara, R. T., & Bassett, D. S. (2024). Structure–function coupling in macroscale human brain networks. *Nature Reviews Neuroscience*, *25*(10), 688-704. https://doi.org/10.1038/s41583-024-00846-6

Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L, … & WU-Minn HCP Consortium. (2013). The minimal preprocessing pipelines for

the Human Connectome Project. Neuroimage, 80, 105-124.
https://doi.org/10.1016/j.neuroimage.2013.04.127

Glasser, M. F., Goyal, M. S., Preuss, T. M., Raichle, M. E., & Van Essen, D. C. (2014).
Trends and properties of human cerebral cortex: correlations with cortical myelin content.
Neuroimage, 93, 165-175. https://doi.org/10.1016/j.neuroimage.2013.03.060

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., ...
& Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. Nature,
536(7615), 171-178. https://doi.org/10.1038/nature18933

Guo, S., Zhang, G., Zeng, X., Xiong, Y., Xu, Y., Cui, Y., ... & Guo, D. (2025). Ten years
of the digital twin brain: Perspectives and challenges. *Europhysics Letters*.
https://doi.org/10.1209/0295-5075/adb3c9

Jeong, S. O., Kang, J., Pae, C., Eo, J., Park, S. M., Son, J., & Park, H. J. (2021).
Empirical Bayes estimation of pairwise maximum entropy model for nonlinear brain state
dynamics. *NeuroImage*, 244, 118618. https://doi.org/10.1016/j.neuroimage.2021.118618

Kandeepan, S., Rudas, J., Gomez, F., Stojanoski, B., Valluri, S., Owen, A. M., ... &
Soddu, A. (2020). Modeling an auditory stimulated brain under altered states of
consciousness using the generalized Ising model. *Neuroimage*, 223, 117367.
https://doi.org/10.1016/j.neuroimage.2020.117367

Kang, J., Jeong, S. O., Pae, C., & Park, H. J. (2021). Bayesian estimation of maximum
entropy model for individualized energy landscape analysis of brain state dynamics.
Human brain mapping, 42(11), 3411-3428. https://doi.org/10.1002/hbm.25442

Kang, J., Lee, H., Yu, S., Lee, M., Kim, H. J., Kwon, R., ... & Yon, D. K. (2024). Effects
and safety of transcranial direct current stimulation on multiple health outcomes: an
umbrella review of randomized clinical trials. *Molecular Psychiatry*, 1-13.
https://doi.org/10.1038/s41380-024-02624-3

Khanra, P., Nakuci, J., Muldoon, S., Watanabe, T., & Masuda, N. (2024). Reliability of
energy landscape analysis of resting-state functional MRI data. *European Journal of
Neuroscience*. https://doi.org/10.1111/ejn.16390

Kristanto, D., Hildebrandt, A., Sommer, W., & Zhou, C. (2023). Cognitive abilities are
associated with specific conjunctions of structural and functional neural subnetworks.
NeuroImage, 279, 120304. https://doi.org/10.1016/j.neuroimage.2023.120304

Liégeois, R., Santos, A., Matta, V., Van De Ville, D., & Sayed, A. H. (2020). Revisiting
correlation-based functional connectivity and its relationship with structural connectivity.
Network Neuroscience, 4(4), 1235-1251. https://doi.org/10.1162/netn_a_00166

Lu, W., Du, X., Wang, J., Zeng, L., Ye, L., Xiang, S., ... & Feng, J. (2024). Simulation and assimilation of the digital human brain. *Nature Computational Science*, 1-9. https://doi.org/10.1038/s43588-024-00731-3

Lynn, C. W., & Bassett, D. S. (2019). The physics of brain network structure, function and control. *Nature Reviews Physics*, *1*(5), 318-332. https://doi.org/10.1038/s42254-019-0040-8

Marinazzo, D., Pellicoro, M., Wu, G., Angelini, L., Cortés, J. M., & Stramaglia, S. (2014). Information transfer and criticality in the Ising model on the human connectome. *PloS One*, 9(4), e93616. https://doi.org/10.1371/journal.pone.0093616

Mišić, B., & Sporns, O. (2016). From regions to connections and networks: new bridges between brain and behavior. *Current opinion in neurobiology*, 40, 1-7. https://doi.org/10.1016/j.conb.2016.05.003

Mueller, S., Wang, D., Fox, M. D., Yeo, B. T., Sepulcre, J., Sabuncu, M. R., ... & Liu, H. (2013). Individual variability in functional connectivity architecture of the human brain. *Neuron*, *77*(3), 586-595. https://doi.org/10.1016/j.neuron.2012.12.028

Nguyen, H. C., Zecchina, R., & Berg, J. (2017). Inverse statistical problems: from the inverse Ising problem to data science. *Advances in Physics*, 66(3), 197-261. https://doi.org/10.1080/00018732.2017.1341604

Palva, J. M., & Palva, S. (2012). Infra-slow fluctuations in electrophysiological recordings, blood-oxygenation-level-dependent signals, and psychophysical time series. *Neuroimage*, 62(4), 2201-2211. https://doi.org/10.1016/j.neuroimage.2012.02.060

Paquola, C., Vos De Wael, R., Wagstyl, K., Bethlehem, R. A., Hong, S. J., Seidlitz, J., ... & Bernhardt, B. C. (2019). Microstructural and functional gradients are increasingly dissociated in transmodal cortices. *PLoS biology*, *17*(5), e3000284. https://doi.org/10.1371/journal.pbio.3000284

Ponce-Alvarez, A., Uhrig, L., Deco, N., Signorelli, C. M., Kringelbach, M. L., Jarraya, B., & Deco, G. (2022). Macroscopic quantities of collective brain activity during wakefulness and anesthesia. *Cerebral cortex*, 32(2), 298-311. https://doi.org/10.1093/cercor/bhab209

Potter, C. C., & Swendsen, R. H. (2013). Guaranteeing total balance in Metropolis algorithm Monte Carlo simulations. *Physica A: Statistical Mechanics and its Applications*, 392(24), 6288-6299. https://doi.org/10.1016/j.physa.2013.08.059

Roudi, Y., Tyrcha, J., & Hertz, J. (2009). Ising model for neural data: model quality and approximate methods for extracting functional connectivity. *Physical Review E—*

Statistical, Nonlinear, and Soft Matter Physics, 79(5), 051915.
https://doi.org/10.1103/PhysRevE.79.051915

Ruffini, G., Damiani, G., Lozano-Soldevilla, D., Deco, N., Rosas, F. E., Kiani, N. A., ... & Deco, G. (2023). LSD-induced increase of Ising temperature and algorithmic complexity of brain dynamics. *PLoS computational biology*, 19(2), e1010811. https://doi.org/10.1371/journal.pcbi.1010811

Stramaglia, S., Pellicoro, M., Angelini, L., Amico, E., Aerts, H., Cortés, J. M., ... & Marinazzo, D. (2017). Ising model with conserved magnetization on the human connectome: Implications on the relation structure-function in wakefulness and anesthesia. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(4). https://doi.org/10.1063/1.4978999

Sultana, O. F., Bandaru, M., Islam, M. A., & Reddy, P. H. (2024). Unraveling the complexity of human brain: Structure, function in healthy and disease states. *Ageing Research Reviews*, *100*, 102414. https://doi.org/10.1016/j.arr.2024.102414

Theis, N., Bahuguna, J., Rubin, J. E., Cape, J., Iyengar, S., & Prasad, K. M. (2024). Diagnostically distinct resting state fMRI energy distributions: A subject-specific maximum entropy modeling study. *bioRxiv*, 2024-01. https://doi.org/10.1101/2024.01.23.576937

Theis, N., Bahuguna, J., Rubin, J. E., Banerjee, S. S., Muldoon, B., & Prasad, K. M. (2025). Energy of Functional Brain States Correlates with Cognition in Adolescent-Onset Schizophrenia and Healthy Persons. *European Journal of Neuroscience*. (Vol. 46, No. 1, p. e70129). Hoboken, USA: John Wiley & Sons, Inc. https://doi.org/10.1002/hbm.70129

Tu, W., Cramer, S. R., & Zhang, N. (2024). Disparity in temporal and spatial relationships between resting-state electrophysiological and fMRI signals. *Research Square*. https://doi.org/10.7554/eLife.95680.3

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. https://doi.org/10.1006/nimg.2001.0978

Venkadesh, S., & Van Horn, J. D. (2021). Integrative models of brain structure and dynamics: concepts, challenges, and methods. *Frontiers in Neuroscience*, *15*, 752332. https://doi.org/10.3389/fnins.2021.752332

Wang, P., Kong, R., Kong, X., Liégeois, R., Orban, C., Deco, G., ... & Thomas Yeo, B. T. (2019). Inversion of a large-scale circuit model reveals a cortical hierarchy in the dynamic resting human brain. *Science advances*, 5(1), eaat7854. https://doi.org/10.1126/sciadv.aat7854

Wang, R., Lin, P., Liu, M., Wu, Y., Zhou, T., & Zhou, C. (2019). Hierarchical connectome modes and critical state jointly maximize human brain functional diversity. *Physical review letters*, *123*(3), 038301. https://doi.org/10.1103/PhysRevLett.123.038301

Wang, R., Liu, M., Cheng, X., Wu, Y., Hildebrandt, A., & Zhou, C. (2021). Segregation, integration, and balance of large-scale resting brain networks configure different cognitive abilities. *Proceedings of the National Academy of Sciences*, 118(23), e2022288118. https://doi.org/10.1073/pnas.2022288118

Watanabe, T., Hirose, S., Wada, H., Imai, Y., Machida, T., Shirouzu, I., ... & Masuda, N. (2013). A pairwise maximum entropy model accurately describes resting-state human brain networks. *Nature communications*, *4*(1), 1370. https://doi.org/10.1038/ncomms2388

WU-Minn, H. C. P. (2017). 1200 subjects data release reference manual. URL https://www. humanconnectome.org, 565, 2.

Zhao, M., Zhou, C., Chen, Y., Hu, B., & Wang, B. H. (2010). Complexity versus modularity and heterogeneity in oscillatory networks: combining segregation and integration in neural systems. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 82(4), 046225. https://doi.org/10.1103/PhysRevE.82.046225

Zimmermann, J., Griffiths, J. D., & McIntosh, A. R. (2018). Unique mapping of structural and functional connectivity on cognition. *Journal of Neuroscience*, 38(45), 9658-9667. https://doi.org/10.1523/JNEUROSCI.0900-18.2018