Title

Bayesian inference of a spectral graph model for brain oscillations

Permalink

https://escholarship.org/uc/item/0933g3vs

Authors

Jin, Huaqing

Verma, Parul

Jiang, Fei

et al.

Publication Date

2023-10-01

DOI

10.1016/j.neuroimage.2023.120278

Peer reviewed

Bayesian inference of a spectral graph model for brain oscillations

Huaqing Jin[a,1], **Parul Verma**[a,1], **Fei Jiang**[b], **Srikantan S Nagarajan**[a], **Ashish Raj**[a,1]

[a]Department of Radiology and Biomedical Imaging University of California San Francisco, San Francisco, CA, USA

[b]Department of Epidemiology and Biostatistics University of California San Francisco, San Francisco, CA, USA

Abstract

The relationship between brain functional connectivity and structural connectivity has caught extensive attention of the neuroscience community, commonly inferred using mathematical modeling. Among many modeling approaches, spectral graph model (SGM) is distinctive as it has a closed-form solution of the wideband frequency spectra of brain oscillations, requiring only global biophysically interpretable parameters. While SGM is parsimonious in parameters, the determination of SGM parameters is non-trivial. Prior works on SGM determine the parameters through a computational intensive annealing algorithm, which only provides a point estimate with no confidence intervals for parameter estimates. To fill this gap, we incorporate the simulation-based inference (SBI) algorithm and develop a Bayesian procedure for inferring the posterior distribution of the SGM parameters. Furthermore, using SBI dramatically reduces the computational burden for inferring the SGM parameters. We evaluate the proposed SBI-SGM framework on the resting-state magnetoencephalography recordings from healthy subjects and show that the proposed procedure has similar performance to the annealing algorithm in recovering power spectra and the spatial distribution of the alpha frequency band. In addition, we also analyze the correlations among the parameters and their uncertainty with the posterior distribution which cannot be done with annealing inference. These analyses provide a richer understanding of the interactions among biophysical parameters of the SGM. In general, the use of simulation-based Bayesian inference enables robust and efficient computations of generative model parameter uncertainties and may pave the way for the use of generative models in clinical translation applications.

Keywords

Bayesian; Connectomes; Magnetoencephalography; Spectral graph theory; Simulation-based inference

1. Introduction

A key endeavor in the field of neuroscience is to uncover the relationship between the brain's complex electrophysiological and functional activity, and its underlying structural wiring contained in white matter fiber projections (Fornito et al., 2015; Suárez et al., 2020). Functional activity between the gray matter regions is estimated with functional magnetic resonance imaging (fMRI), electroencephalography (EEG), and magnetoencephalography (MEG), while the structural wiring is assessed using diffusion tensor imaging from MRI. The brain structure-function relationship is then investigated using various data-driven and mathematical modeling-based techniques, assuming structural connectivity as a graph with different brain regions as graph nodes connected to each other via edges that are informed by the white matter fiber projections.

While both data-driven (Strogatz, 2001; Buckner et al., 2005; Achard et al., 2006; Bassett and Bullmore, 2006; Chatterjee and Sinha, 2007; He et al., 2008; Ghosh et al., 2008; Bullmore and Sporns, 2009; Rubinov et al., 2009; van den Heuvel et al., 2009; Bassett and Bullmore, 2009; Abdelnour et al., 2014; Hermundstad et al., 2013; Park and Friston, 2013; Abdelnour et al., 2018) as well as modeling techniques (Wilson and Cowan, 1973; David and Friston, 2003; Destexhe and Sejnowski, 2009; El Boustani and Destexhe, 2009; Honey et al., 2009; Spiegler and Jirsa, 2013; Cabral et al., 2014; Muldoon et al., 2016; Siettos and Starke, 2016; Cabral et al., 2017; Breakspear, 2017) have been employed to uncover the brain structure-function relationship, mathematical models additionally provide insights into the underlying biophysics of brain activity. After fitting the model to empirical fMRI, EEG, and MEG data, the inferred model parameters can serve as biophysically interpretable markers of disease and brain states (Honey and Sporns, 2008; Alstott et al., 2009; Haan et al., 2012; Yang et al., 2016; Zimmermann et al., 2018; Singh et al., 2020). For example, Zimmermann et al. Zimmermann et al. (2018) demonstrated that the model parameters can predict cognition. However, the practical impact of model-based biomarkers of pathophysiology is hampered by two key challenges, described below.

Lack of interval estimations and posterior probabilities.

An important goal of practical model fitting is to quantify how well a model's parameters explain empirical neuroimaging data, and how confidently those estimates can be obtained. It is, therefore, necessary to capture their variability and find out all possible parameter settings compatible with the observed phenomena (Gonçalves et al., 2020). Bayesian inference is the established approach for achieving these goals, by making available the posterior distribution of parameters given the observations. Posterior distribution in turn provides rich information about how model parameters interact together, and quantifies the uncertainty of the model output — potentially critical for obtaining computational biomarkers in disease. Unfortunately, Bayesian inference methods have been proven to be

quite challenging for most current computational models of brain activity (Gonçalves et al., 2020; Cranmer et al., 2020; Van Geit et al., 2016).

Tractability of model inference.

We identify three issues limiting the tractability of Bayesian model evidence in the field. First, powerful sampling methods like Markov Chain Monte Carlo (MCMC) require extremely large samples, numbering in the hundreds of thousands. Most current models, like the coupled neural mass models (NMMs), are evaluated via time-consuming numerical integration techniques, which in turn impose a prohibitive computational burden on any sampling technique. Second, coupled NMMs involve large parameter spaces, i.e. number of internal parameters that must be jointly inferred, making full Bayesian inference impractical. Third, due to inherent non-linearity, the theoretical posterior density in even the simplest computational models is so convoluted, non-smooth, and non-convex that conventional optimization or MCMC sampling techniques encounter huge challenges. Many of these issues are highlighted in previous studies (Hartoyo et al., 2019; Raj et al., 2022; Xie et al., 2019), and together they have ensured that hardly any Bayesian inference is performed in these settings.

In this paper, we present a novel way for Bayesian inference of computational models of neural activity, focusing specifically on the recently proposed spectral graph model (SGM), a linear biophysical generative model that can accurately capture the steady state wideband power spectral density (PSD) as well as the spatial distribution of the alpha band obtained from MEG (Raj et al., 2020). We choose the SGM for the following reasons:

1. SGM involves a parsimonious set of global biophysically interpretable parameters; in our previous paper, we demonstrated that only 7 global, spatially-invariant parameters, each having distinct biophysical meaning, were sufficient to accurately capture empirical MEG PSD (Verma et al., 2022a,b). This may be compared against previous models that have typically required substantially more spatially-varying parameters.

2. SGM explicitly estimates regional PSD and therefore can directly fit the frequency PSD obtained from MEG/EEG. Other models typically provide time-domain simulations only, and their spectral content is usually not a target of model fitting.

3. SGM is extremely fast to evaluate since its solution can be obtained in a closed-form in the frequency domain. Other models typically require lengthy time-domain simulations, which can be impractical in MCMC or other sampling techniques.

As a result of its linearity and closed-form evaluation without the need for long simulations, SGM parameter inference is far more tractable compared to non-linear neural mass models — where identifiability of model parameter is not guaranteed (Hartoyo et al., 2019; Raj et al., 2022; Xie et al., 2019).

In our prior works we had estimated SGM parameters using a global optimization algorithm, the dual annealing method, as point estimates (Raj et al., 2020; Verma et al., 2022a,b; Raj

et al., 2022), as a preferred alternative over continuous gradient descent-based minimization. However, since the objective function of estimating SGM parameter (Raj et al., 2020; Verma et al., 2022a) is non-convex, the annealing approach does not guarantee a global optimum. Moreover, a single point estimate from the annealing method is far from enough to uncover the underlying entire range and behavior of biophysical processes and to lead to new insights. On the contrary, Bayesian method allows the estimation of the full posterior distribution of the SGM parameters, which is necessary for biological interpretation. As a result, Bayesian method is more suitable for inferring the SGM parameters. However, the conventional Bayesian inference is challenging due to the fact that the theoretical posterior density of SGM parameters may be rather complicated which causes difficulty in sampling.

To circumvent the computational difficulty, we propose a novel method to perform Bayesian inference of the SGM parameters. The method approximates the posterior density of the SGM parameters by using a neural network model, which is trained through a simulation-based inference (SBI) framework (Cranmer et al., 2020). Thus the method is referred to as SBI-SGM. Our main contribution is to show that this custom combination of SGM with SBI is exquisitely well-matched for estimating posterior distribution of generative model parameters.

This provides a far more appealing practical utility, which may be exploited in future clinical applications. Given its speed, this tool can be used to quickly infer posteriors of model parameters for a large number of subjects which can subsequently be used to identify parsimonious markers of disease and brain states. In particular, SBI-SGM can capture the uncertainty of the estimates via providing the credible intervals, which are instrumental in deriving statistically robust conclusions. Once extensive validation on larger samples have been demonstrated, our approach can be a meaningful step in alleviating the critical issue of lack of large sets of empirical data.

Applying the SBI tool to the SGM, we demonstrate that the posteriors can accurately capture the empirical spatial distribution of alpha frequency band and PSD in MEG, and the inference of posteriors is substantially faster than the point estimate inference algorithm used in prior works. This combination of a fast and parsimonious forward model (SGM) with a fast neural network for posterior inference (SBI) is not currently available in the field of structure-function relationship, and could constitute a critical advance in the applicability of computational models to practical scenarios.

2. Methods

2.1. Dataset

We study the resting-state MEG data obtained from 36 healthy subjects (23 males, 13 females; 26 left-handed, 10 right-handed; mean age 21.75 years, age range 7–51 years) as also reported in Raj et al.'s study (Raj et al., 2020; Xie et al., 2020). All study procedures were approved by the institutional review board at the University of California at San Francisco and were in accordance with the ethics standards of the Helsinki Declaration of 1975 as revised in 2008. MRI followed by tractography was used to generate the connectivity and distance matrices. The publicly available dataset consisted of processed

connectivity and distance matrices, and PSD for every subject. The data collection procedure was described in Raj et al. (2020), Verma et al. (2022a) as well as summarized below.

MRI.—A 3 T TIM Trio MR scanner (Siemens, Erlangen, Germany) was used to perform MRI using a 32-channel phased-array radiofrequency head coil. High-resolution MRI of each subject's brain was collected using an axial 3D magnetization prepared rapid-acquisition gradient-echo T1-weighted sequence (echo time [TE] = 1.64 ms, repetition time [TR] = 2530 ms, TI = 1200 ms, flip angle of 7°) with a 256-mm field of view, and 160 1.0-mm contiguous partitions at a 256×256 matrix. Whole-brain diffusion-weighted images were collected at $b = 1,000$ s/mm^2 with 30 directions using 2-mm voxel resolution in-plane and through-plane.

Region parcellations.—The T1-weighted images were parcellated into 68 cortical regions and 18 subcortical regions using the Desikan–Killiany atlas available in the FreeSurfer software (Fischl et al., 2002). To do this, the subject-specific T1-weighted images were back-projected to the atlas using affine registration, as described in the previous studies (Abdelnour et al., 2014; Owen et al., 2013).

Structural connectivity networks.—Different structural connectivity networks were reconstructed with the same Desikan–Killiany parcellations. Firstly, openly available diffusion MRI data were obtained from the MGH-USC Human Connectome Project to create an average template connectome. As in previous studies (Abdelnour et al., 2014; Owen et al., 2013), subject-specific structural connectivity was computed on diffusion MRI data: *Bedpostx* was used to determine the orientation of brain fibers in conjunction with *FLIRT*, as implemented in the *FSL software* (Jenkinson et al., 2012). In order to determine the elements of the adjacency matrix, tractography was performed using *probtrackx2*. 4000 streamlines were initiated from each seed voxel corresponding to a cortical or subcortical gray matter structure and how many of these streamlines reached a target gray matter structure was tracked. The weighted connection between the two structures $c_{i,j}$, was defined as the number of streamlines initiated by voxels in region i that reach any voxel within region j, normalized by the sum of the source and target region volumes ($c_{i,j} = \frac{streamlines}{v_i + v_j}$).

Afterward, connection strengths were averaged between both directions ($c_{i,j}$ and $c_{j,i}$) to form undirected edges. To determine the geographic location of an edge, the top 95% of nonzero voxels by streamline count were computed for both edge directions. The consensus edge was defined as the union between both post-threshold sets.

MEG data.—MEG recordings were acquired at UCSF using a 275-channel CTF Omega 2000 whole-head MEG system from VSM MedTech (Coquitlam, BC, Canada). All subjects were instructed to keep their eyes closed for 5 min while their MEGs were recorded at a sampling frequency of 1200 Hz.

MEG processing and source reconstruction.—MEG recordings were down-sampled from 1200 Hz to 600 Hz, then digitally filtered to remove DC offset and any other noisy artifact outside of the 1 to 160 Hz bandpass range. Since MEG data are in sensor space, meaning they represent the signal observable from sensors placed outside the head, this

data needed to be "inverted" in order to infer the neuronal activity that had generated the observed signal by solving the so-called inverse problem. Adaptive spatial filtering algorithms were used from the NUTMEG software tool written in house (Dalal et al., 2004). To prepare for source localization, all MEG sensor locations were co-registered to each subject's anatomical MRI scans. The lead field (forward model) for each subject was calculated in NUTMEG using a multiple local-spheres head model (three-orientation lead field) and an 8 mm voxel grid which generated more than 5000 dipole sources, all sources were normalized to have a norm of 1. Finally, the MEG recordings were projected into the source space using a beamformer spatial filter. Only the sources belonging to the 68 cortical regions were selected to be averaged around the centroid. All dipole sources were labeled based on the Desikan–Killiany parcellations, then sources within a 20 mm radial distance to the centroid of each brain region were extracted, and the average time course of each region's extracted sources served as empirical resting-state data for our proposed model. MEG recordings were bandpass filtered between 2 to 45 Hz using firls in MATLAB (Anon, 2020) and the static frequency PSD was generated for every region of interest using the pmtm algorithm in MATLAB (Anon, 2020).

2.2. Spectral graph model

SGM is a hierarchical, linear, analytic model of brain oscillations, which has a closed-form solution in the Fourier frequency domain via the eigen-decomposition of a graph Laplacian (Raj et al., 2020; Verma et al., 2022a,b). A typical SGM has two model layers, a mesoscopic layer for the local neuronal subpopulations of every brain region and a macroscopic layer for the long-range excitatory neuronal subpopulations of the whole brain. SGM is briefly described below, and detailed illustrations can be found in the supplementary document and in prior publications (Raj et al., 2020; Verma et al., 2022a).

SGM is characterized by seven identifiable parameters $\mathbf{s} = (\tau_e, \tau_i, \alpha, v, g_{ei}, g_{ii}, \tau_G)^{\mathrm{T}}$, which include the excitatory and inhibitory time constants τ_e, τ_i and neural gains g_{ei} and g_{ii} at the mesoscopic level, and long-range excitatory time constant τ_G, coupling constant α, speed v at the macroscopic level. Given the signals $[x_1(t), ..., x_N(t)]^{\mathrm{T}}$ with N regions of interest (ROIs) in the time domain, the closed-form solution of SGM is obtained in the Fourier domain:

$$\mathbf{X}(\mathbf{s}, \omega) = [\mathscr{F}(x_1(t)), ..., \mathscr{F}(x_N(t))]^{\mathrm{T}} = \left(j\omega\mathbf{I} + \frac{1}{\tau_G}F_G(\tau_G; \omega)\mathscr{L}(\alpha, v; \omega)\right)^{-1} H_{\mathrm{local}}(\tau_e, \tau_i, g_{ei}, g_{ii}; \omega)\mathbf{P}(\omega),$$

(1)

where ω is the angular frequency, $\mathbf{X}(\mathbf{s}, \omega)$ is a vector of the Fourier transformation, or equivalently the PSD, of the macroscopic signal over all brain regions of interest at frequency ω, $\mathscr{F}$ is the Fourier transformation, $\mathscr{L}$ is the complex Laplacian, H_{local} is the mesoscopic model's transfer function, $\mathbf{P}(\omega)$ is the input noise spectrum, and $F_G(\omega)$ is the Fourier transform of a Gamma-shaped neural response function, given as $\frac{1/\tau_G^2}{(j\omega + 1/\tau_G)^2}$. This response function is governed by the characteristic long-range excitatory time constant τ_G, and the function is intended to serve as a lumped model of various processes, including

dendritic arborization, somatic conductance, synaptic capacitance, etc (Raj et al., 2020). To facilitate the illustration of SGM, a schematic is presented in the top panel of Fig. 1.

2.3. Simulation-based inference for SGM

SBI is a powerful tool for the inference of large complex statistical models that have been extensively applied in many areas of science and engineering (Papamakarios and Murray, 2016; Greenberg et al., 2019; Durkan et al., 2019). We adapt the SBI method for SGM parameter estimation and inference. Let $\mathbf{X}(\mathbf{s}, \Omega) = \{\mathbf{X}(\mathbf{s}, \omega)\}_{\omega \in \Omega}$ be the model output PSD in dB scale (Persson and Björkman, 1988) where Ω is the set of the frequency points we used and it contains 40 equally spaced frequencies in the range 2–45 Hz in the manuscript. $G\{\mathbf{X}(\mathbf{s}, \Omega)\}$ is a monotonic transformation that standardizes the PSD across the frequency into a z-score; standardizes the regional distribution of alpha band power (i.e. summation of PSD from 8–12 Hz); and finally concatenates both into a single vector. Here and throughout the text, we present the PSD in dB scale. In our SBI-SGM framework, we assume the data model is

$$\mathbf{y} = G\{\mathbf{X}(\mathbf{s}, \Omega)\} + \epsilon,$$

(2)

where $\epsilon \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$ is additive i.i.d. noise with standard deviation (SD) σ and $\mathbf{I}$ is an identity matrix.

The random noise in (2) captures the model misfitting error in SGM model. We adopt the Gaussian noise for convenience; indeed, SBI-SGM demonstrates resilience concerning the choice of the noise distribution, provided that the standard deviation is well controlled, as delineated in Section S.2 of the supplementary material. Without this random noise, the target posterior density is discontinuous, which is difficult to estimate due to the well-known Gibbs phenomenon (Llanas et al., 2008). Adding this random noise to the model results in a smooth posterior distribution of SGM parameters, which can be accurately approximated by a neural network (Leshno et al., 1993).

Similarly, we define the observed sample $\mathbf{y}_o = G\{\mathbf{X}_{\text{MEG}}(\Omega)\}$ where $\mathbf{X}_{\text{MEG}}(\Omega)$ is the observed MEG PSD.

Since the parameters in the SGM model are assumed to be bounded to satisfy biological constraints, bounded priors are typically adopted for them, which causes difficulties for posterior sampling with SBI (Deistler et al., 2022). To address this issue, we re-parameterize the parameters so that the posterior sampling can be performed on the real line. More specifically, let $\mathscr{H}(x)$ be a scaled logit transformation function (Hilbe, 2009) defined as

$$\mathscr{H}(x) \equiv 10 \times \log\left\{\frac{(x - x_l)/(x_u - x_l)}{1 - (x - x_l)/(x_u - x_l)}\right\}, \text{ for } x \in [x_l, x_u],$$

where x_l and x_u are lower and upper bounds of variable x, respectively. Slightly abusing the notation, let $\theta = \mathscr{H}(\mathbf{s})$, where $\mathscr{H}(\mathbf{s})$ represents the values of function $\mathscr{H}$ applied on each element of $\mathbf{s}$.

Under a Bayesian framework, we are interested in the posterior distribution of θ given $\mathbf{y}$, particularly the Bayesian credible interval of θ, which captures the uncertainty of the SGM parameters. To obtain the credible interval, we estimate the posterior distribution of θ through the SBI procedure (Cranmer et al., 2020). The density of $\mathbf{y}$ is denoted by $p(\mathbf{y} \mid \theta)$ following (2), which is a multivariate Gaussian density function. We impose a multivariate Gaussian prior $\pi(\theta)$ on θ. The posterior density is $q_{\Phi}(\theta \mid \mathbf{y}) \propto \pi(\theta)p(\mathbf{y} \mid \theta)$, where Φ is the unknown parameters that determine the posterior distribution. Instead of obtaining the posterior density for the SGM parameters s directly, we first derive the posterior density for θ, which results in the target posterior distribution through a Jacobian transformation (Henderson and Searle, 1979).

We use a deep learning architecture, namely neural spline flow (NSF) (Durkan et al., 2019), to model the functional form of q_{Φ}, where Φ is the parameter in the deep learning network. The dimension of Φ increases with the number of network layers in NSF, and when the dimension of Φ approaches infinity, q_{Φ} approaches the true posterior distribution. When the deep learning architecture is given, Φ is the only unknown parameter in $q_{\Phi}(\theta \mid \mathbf{y})$. Hence estimating the posterior density is equivalent to estimating Φ. Now note that the true posterior distribution maximizes $\mathbf{E}[\log\{q_{\Phi}(\theta \mid \mathbf{y})\}]$, where the expectation is taken with respect to $\mathbf{y}$ and θ, we propose to obtain an estimator for Φ through maximizing the empirical version of $\mathbf{E}[\log\{q_{\Phi}(\theta \mid \mathbf{y})\}]$, that is $\frac{1}{M}\sum_{m=1}^{M} \log\{q_{\Phi}(\theta_m \mid \mathbf{y}_m)\}$, where the samples $\mathbf{y}_m$ and θ_m, $m = 1, \ldots, M$ are the simulated realizations of $\mathbf{y}$ and θ based on $p(\mathbf{y} \mid \theta)$ and $\pi(\theta)$, respectively.

To obtain the posterior density for s given the observed sampled from the empirical PSD of MEG data $\mathbf{y}_o = G\{\mathbf{X}_{\mathrm{MEG}}(\Omega)\}$, we can feed $\mathbf{y}_o$ in the neural network and obtain the estimated posterior distribution $q_{\Phi}(\theta \mid \mathbf{y}_o)$ with the estimated parameter $\widehat{\Phi}$. The target posterior distribution of $\mathbf{s}$ is $q_{\widehat{\Phi}} \times |\det(\mathbf{J})|$ where $\mathbf{J}$ is the Jacobian matrix, i.e., $\mathbf{J} = \partial\theta / \partial\mathbf{s}$. We illustrate the details of obtaining the posterior distribution of θ in Algorithm 1, which contains a simulation step and an optimization step.

Algorithm 1

Posterior estimation with re-parameterization

Require: A multivariate Gaussian prior $\pi(\theta) \sim N(0,\ 100\mathbf{I})$ a likelihood $p(\mathbf{y} \mid \theta)$, an observation $\mathbf{y}_o$.

Simulation:

for $m = 1, \ldots, M$ **do**

Sample $\theta_m \sim \pi(\theta)$

Sample $\mathbf{y}_m \sim p(\mathbf{y} \mid \theta_m)$

end for

Optimization:

$$\widehat{\boldsymbol{\Phi}} \leftarrow \underset{\boldsymbol{\Phi}}{\operatorname{argmin}} - \frac{1}{M} \sum_{m=1}^{M} \log\{q_{\boldsymbol{\Phi}}(\theta_m \mid \mathbf{y}_m)\}$$

return $q_{\widehat{\boldsymbol{\Phi}}}(\theta \mid \mathbf{y}_o)$ as the estimate of the posterior distribution of θ. The posterior distribution of $\mathbf{s}$ is $q_{\widehat{\boldsymbol{\Phi}}} \times |\det(\mathbf{J})|$, where $\mathbf{J}$ is the Jacobian matrix, i.e., $\mathbf{J} = \partial\boldsymbol{\theta}/\partial\mathbf{s}$.

2.4. Implementation details

For each subject, we use their MEG data from 68 cortical regions according to the Desikan–Killiany atlas (Desikan et al., 2006) to obtain the posterior samples of the SGM parameters s. Under this atlas, we obtain a 68-region $X(s, \omega)$ at frequency ω, and the dimension of $G\{X(s, \Omega)\}$ is $68 \times 40 + 68 = 2788$ which is comprised of both the PSD feature and the spatial feature in the alpha band.

We implement SBI-SGM using the sbi package in Python (https://www.mackelab.org/sbi/) (Tejero-Cantero et al., 2020), where the hyperparameters in the original SBI algorithms are adopted as the default values provided in the package. We discuss the choice of the standard deviation of the noise σ and the number of simulation samples in <u>Simulation</u> step in Algorithm 1 in the next section.

In SBI-SGM, we adopt an average template structural connectome created via openly available diffusion MRI data obtained from the MGH-USC human connectome project (HCP) for training a universal posterior mapping from observation to the posterior distribution using Algorithm 1. After obtaining a trained posterior density for each observed $\mathbf{y}_o$, we draw a posterior sample of SGM parameters, denoted by $\widetilde{\mathbf{s}}$. We then obtain $\mathbf{X}(\widetilde{\mathbf{s}}, \Omega)$ using (1) and the individual structural connectome. Finally, we construct the standardized PSD and spatial distribution of the alpha band PSD as a G transformation of $\mathbf{X}(\widetilde{\mathbf{s}}, \Omega)$, where the function G is defined in (2). We perform this posterior sampling process 1000 times to obtain a set of posterior samples of the standardized PSD and spatial distribution of the alpha band PSD for each observed $\mathbf{y}_o$. The pipeline of SBI inference for SGM is presented in the bottom panel of Fig. 1.

We compare the performance of SBI-SGM with the performance of the annealing SGM approach (Raj et al., 2020; Verma et al., 2022a,b), namely Ann-SGM, on our MEG data. The details of the annealing implementation can be found in Verma et al. (2022a). SGM model assumes the parameters have finite supports as the ones listed in Table 1. In Ann-SGM, three different bounds are evaluated for g_{ei} and g_{ii} sequentially, and the largest bounds that satisfy the stability condition defined by Verma et al. (2022b) are chosen in the subsequent estimation. In SBI-SGM, the largest bounds in Table 1 are adopted for parameters (g_{ei}, g_{ii}) and we only retain the posterior samples of the SGM parameters within the stability boundary defined in Verma et al. (2022b). For the other five parameters, SBI-SGM uses the same bounds as Ann-SGM does.

3.   Results

3.1.   Adding random noise to the SGM improves the reconstructing accuracy of the PSD

For each subject, we obtain the reconstructed PSD by taking the mean of the posterior samples of the PSD. We then study how the change in noise variation affects the performance of SBI-SGM in reconstructing the observed PSD. We compare the median Pearson's correlation between the reconstructed PSD and the observed PSD from MEG. Specifically, for each ROI, we calculate the correlation between the reconstructed PSD and observed PSD from MEG. We then average the correlations over all ROIs and obtain the median of this average correlation over 36 subjects. In this study, the number of simulations samples is fixed at 100,000 in the <u>Simulation</u> step in Algorithm 1 and the standard deviation of ϵ varies from 0 to 3.2. We report the mean results over 10 repetitions. Note that when $\sigma = 0$, there is no random noise added.

Fig. 2A shows that compared to the model without random error (when the $\sigma = 0$), adding random noise in (2) significantly reduces the reconstruction errors. This result is consistent with our theoretical conclusion that adding random noise results in a smooth posterior density which can be accurately approximated by a neural network. The Pearson's correlation between the reconstructed and the observed PSD increases when $\sigma < 1.6$ and starts to decrease after σ reaches 1.6, when the signal-to-noise ratio is not sufficiently large for the SBI-SGM to recover the observed PSD. In practice, we suggest choosing σ in [0.8, 2.0], which yields satisfactory performance with over 0.9 correlation between reconstructed and observed PSD. For all the following experiments, we fix $\sigma = 1.6$.

3.2.   Increasing the number of simulation samples improves the SBI-SGM fit

We also investigate how the performance of SBI-SGM changes with the number of simulations in <u>Simulation</u> step in Algorithm 1. Fig. 2B shows that a larger number of simulations yields a better SBI-SGM fit with a higher correlation between reconstructed and observed PSD. As indicated by the right panel of Fig. 2B, the changes in Pearson's correlation are not very notable and it varies from 0.887 to 0.906. However, when considering the standardized PSD curves for different numbers of samples, a clear trend is observed that typically larger sample size leads to a better fit visually. It is also worth noting that the performance of SBI-SGM is stable after the number of simulations reaches 100,000. Therefore, we choose 100,000 simulations in the <u>Simulation</u> step in Algorithm 1 in the subsequent analyses.

3.3.   Results from two representative MEG data

We show the results from two representative subjects whose Pearson's correlations between reconstructed PSD and the observed PSD are the top two closest to the median correlation across 36 subjects in one experiment. To make it representative, we repeat the SBI-SGM procedure 10 times and choose the experiment which yields an overall correlation closest to the mean level in the 10 repetitions for the analysis.

The posterior samples of the seven parameters as well as the PSD for two subjects are displayed in Figs. 3A and B. In the left panels, we compare the posterior density of

SBI-SGM with the point estimate from Ann-SGM. We can observe multiple modes from the posterior densities. For τ_e, v, and g_{ii}, the point estimates from Ann-SGM are close to one of the modes of the posterior distributions, while the estimates of the rest of the parameters from Ann-SGM are far away from the posterior modes from SBI-SGM.

The middle panels of Fig. 3 shows the posterior mean and 95% credible interval of the PSD from SBI-SGM. In each subject, the 95% credible interval covers the observed PSD at low frequencies (lower than 20 Hz), which is consistent with the fact that SGM can successfully recover the low-frequency PSD (Verma et al., 2022a).

Moreover, we vectorize the reconstructed and observed PSDs and map them onto a 5-dimensional manifold using the uniform manifold approximation and projection (UMAP) method proposed by McInnes et al. (2018). As shown in the right panels of Fig. 3, the projection of the observed PSD falls within the support of the projection of the reconstructed PSDs in the manifold, which further validates our Bayesian inference (Gabry et al., 2017).

3.4. Cohort level analysis of MEG datasets

In SBI, the variability of the posterior distribution exists due to the randomness of the simulated samples in the Simulation step and the randomness in the posterior sampling procedure using the trained posterior distribution. We evaluate the robustness of SBI-SGM in 10 repetitions. In Fig. 4A, we show the median of the reconstructed PSDs over 36 subjects for each repetition, the PSD Pearson's correlation is changed between [0.905,0.907] (shown in the caption). The results indicate that SBI-SGM is robust throughout the repetitions. We choose an experiment that yields a correlation closest to the mean level in the 10 repetitions in the subsequent analyses.

We analyze the posterior samples of the SGM parameters from SBI-SGM across 36 MEG data. We first study the pair-wise correlation between the SGM parameters using the partial correlation method (Marrelec et al., 2006), which examines the correlation between any given two parameters after removing the effect from other parameters, Fig. 4B shows the pair-wise partial correlation averaged over 36 subjects. As shown in Fig. 4B, speed v has no correlation with the other parameters. The two time constants τ_e, τ_i have weak positive correlation, and the graph time constant τ_G shows moderate negative correlation with the exhibitory time constant τ_e and small positive correlation with the inhibitory time constant τ_i. Fig. 4C shows the distribution of the pooled posterior samples of the SGM parameters over 36 MEG data. Among the seven SGM parameters, the posterior distributions of τ_i, τ_G are highly concentrated, which indicates their variabilities across different subjects are small. The histogram of the inhibitory time constant τ_i presents a second peak around 0.15s. The speed v has the highest density round 15 m/s.

We further investigate whether the SGM model in (1) can generate the observed PSDs. We generate 1000 SGM parameters from the prior distribution of s, and obtain simulated PSDs through (1). We then compare the simulated PSDs with the observed ones. To facilitate the visualization, we vectorize the simulated samples of PSDs and observed PSDs and utilize the UMAP method to project them to a 2-dimensional embedding manifold. As shown in

Fig. 5, in the embedded manifold, all the observed projections fall within the projections of the simulated samples, which indicates that the SGM model captures the generating mechanism of the observed PSD, and therefore is a reasonable model of the data.

3.5. SBI-SGM and Ann-SGM comparision

We compare SBI-SGM with Ann-SGM. In Fig. 6A, we show the correlations between the reconstructed and observed PSDs and the correlations between reconstructed and observed spatial distributions of the alpha band PSD resulting from SBI-SGM and Ann-SGM. We also perform statistical tests on the difference between the results from the two inference methods. Specifically, we calculate the Pearson's correlations for each ROI between the reconstructed and observed PSDs and take the average across ROIs. Furthermore, we obtain the spatial correlation as the inner product between the reconstructed and observed spatial distribution of the alpha band PSD weighted by $\mathbf{D} + w\mathbf{I}$ where $\mathbf{D}$ is the row degree normalized structural connectivity matrix, $\mathbf{I}$ is the identity matrix, w is an empirical weight, and we adopt $w = 10$ as suggested by Raj et al. (2022).

As shown in Fig. 6A, SBI-SGM gives similar average correlation and spatial correlation as Ann-SGM does with insignificant p-values from two-sample t-tests. In Fig. 6B, we observe very similar spatial distributions of the alpha band PSD from SBI-SGM and Ann-SGM, and both of them are similar to the observed one from MEG data. In conclusion, SBI-SGM has a similar performance as Ann-SGM on recovering observed PSD and spatial distribution from the alpha band.

3.6. Computational cost of SBI-SGM

While SBI-SGM entails training a neural network, it remains considerably faster than Ann-SGM. In our MEG data with 36 subjects, Ann-SGM takes about 8 h to estimate the SGM parameters for each subject and it necessitates repetition for every subject. On the same machine, SBI-SGM takes approximately 2 h to accomplish the Bayesian inference on SGM parameters for all subjects when trained with 100,000 simulation samples.

In fact, the computational bottleneck for SBI-SGM lies in the generation of the simulation samples, which can effectively leverage the parallel computing. If the simulation samples are produced, the training process only takes 0.34 h with 100,000 samples and the inference time is negligible. Moreover, using a template connectome, as in our case, allows the computational cost of SBI-SGM to be amortized, i.e., once trained, the model can be applied to all subjects without modification. These features render SBI-SGM more desirable in practice compared with other methods, such as Ann-SGM and MCMC (Sengupta et al., 2015, 2016), which require to run the estimation procedure for each subject and are not easily adaptable to parallel computing.

The primary computational burden in generating simulation samples stems from the eigen-decomposition of the Laplacian matrix whose computational complexity is $O(N^3)$, where N is the number of ROIs. During training or inference, the number of ROIs only affects input dimensionality, and exhibits linear complexity with respect to N. Consequently, the total complexity in terms of N is $O(N^3)$. In widely-used atlases (Desikan et al., 2006; Fan

et al., 2016), the number of ROIs remains below 250. Hence, despite the $O\left(N^3\right)$ complexity, applying SBI-SGM to other atlases does not result in an excessive computational burden.

4. Discussion

Models with complex and stochastic simulators have been extensively applied in many areas of science and engineering (Lueckmann et al., 2021). In neuroscience, such computational models are typically built via incorporating biological mechanisms and hypothetical intuitions to explain the observed phenomena inferred from the neuroimaging data (Gonçalves et al., 2020). These models involve several free parameters that are required to be compatible with the observed phenomena. Due to the complexity of neural models and neural data, the determination of the free parameters generally relies on computation-intensive optimization routines like grid search (Tomm et al., 2011), genetic algorithm (Van Geit et al., 2016) or simulated annealing (Raj et al., 2020; Verma et al., 2022a).

However, these algorithms are far from meeting the needs of the neuroscience community, as they can only provide a single point estimate of the free parameters, and make it difficult to incorporate prior knowledge about related neural processes. In neural models, it is always desirable to find out not only the best, but all parameter settings compatible with the observed data. The variability of the parameters under the observation can provide more insights into the neural models and processes (Gonçalves et al., 2020; Alonso and Marder, 2019). Moreover, neural model parameters, e.g. of SGM, typically have biological meaning, hence their inference must accommodate the underlying biological mechanisms and their constraints, in order to avoid unreasonable solutions. Using the prior knowledge of these biological quantities can not only increase the optimization efficiency but also robustify the inferred models. Most importantly, the practical applicability of model fitting demands a solid assessment of the confidence bounds and variability associated with fitted parameters — something quite lacking in current methods. Due to these reasons, full Bayesian inference of posteriors is preferable to point estimates. However, the intricacy of the neural models typically results in intractable or complicated likelihood which makes the likelihood-based inference inaccessible.

Luckily, the SBI approach fills this gap by bypassing the evaluation of the likelihood function and giving the posterior samples directly. The results presented in this study have highlighted the key ways in which the proposed combination of SGM and SBI is exquisitely well-suited to the task of model inference of neural systems. First, the parsimony of SGM obviates a key weakness of SBI, which typically prefers to infer a small set of parameters (Cranmer et al., 2020; Durkan et al., 2019). For this reason, SBI may be challenging for coupled non-linear models such as NMMs and the Virtual Brain (Ritter et al., 2013) which consist of a potentially large set of parameters. Second, SBI requires a large number of forward-model evaluations to generate enough simulation samples for training, which would render large coupled NMMs unfeasible (Cabral et al., 2014; Sanz-Leon et al., 2015; Cabral et al., 2017; Siettos and Starke, 2016), but this is far less problematic for SGM due to its fast forward evaluation. Third, SBI requires far fewer empirical samples compared to simulation samples (Durkan et al., 2019; Cranmer et al., 2020), which is an important consideration in real data-poor medical settings. Lastly, while the training of the neural network requires

a high upfront cost involving numerous simulations, the trained SBI model can be applied almost instantaneously to new empirical data directly, which enhances the practical utility and amortizes the computational cost of fitting by front-loading the simulation effort.

We were able to show that the SBI-SGM framework gives speedy estimates of the full posterior distribution, achievable in a matter of seconds per subject. Using these posteriors, point estimates, e.g. mean or mode of the posterior, can be quickly produced, which we showed has comparable performance to prior point estimation methods like dual annealing, at a fraction of the computation time. Lastly, our posterior analysis showed that the model parameters were generally weakly correlated, implying that all of them are required to obtain model outputs that match the spectral and spatial patterns obtained from empirical MEG. This is a crucial finding since it suggests that we can identify unique markers of diseases and brain states in the form of inferred SGM parameters.

4.1. Relationship to previous works

For models closely related to SGM, such as the non-linear neural mass models or the dynamic causal models (DCM), Bayesian inferencing algorithms such as variational Bayes have been used previously. DCM employs variation Bayes to obtain effective functional connectivity (Kiebel et al., 2008; Pinotsis et al., 2012). A key difference is that DCM is primarily used to obtain effective connectivity from smaller networks and that these connectivities are obtained from second-order statistics such as cross-spectra using spectral DCM (Pereira et al., 2021). In contrast, SGM directly computes the PSD rather than individual elements of the second-order effective connectivity matrix. SGM instead employs an explicit structure-based model, where the inter-regional connectivity comes directly from the measured structural connectome. In this manner, SGM is better suited for SBI than DCM, since the latter would be required to infer an entire matrix of effective connectivities, in addition to other regional or global parameters.

The key challenges with inferring parameters of coupled non-linear neural mass models are that they require time-consuming simulations. These models exhibit bifurcations yielding discontinuities in the model solutions (Raj et al., 2022), and parameter identifiability is not guaranteed (Hartoyo et al., 2019). These challenges have been discussed in detail elsewhere (Xie et al., 2019; Raj et al., 2022). SGM overcomes these challenges by providing a closed-form solution that can be simulated within seconds, and by consisting of only a parsimonious set of 7 global model parameters.

Parsimony is an essential feature and motivation of our study, and indeed the sufficiency of only 7 parameters in SGM to explain whole brain activity is a powerful, if surprising, concept. This aspect was dealt with extensively in prior publications on the spectral graph model (Raj et al., 2020; Verma et al., 2022a; Raj et al., 2022). Here we offer some pertinent comments. It should be noted that the SGM, and other graph models like it, seeks a *minimal* model of brain activity, with highly circumscribed targets. It does not seek to model the full complexity of brain activity, but only its power spectrum. In this sense, it assumes the activity is strictly stationary. Most existing analyses of MEG activity already rely on stationary features like functional or effective connectivity, which represent the vast majority of use-cases.

Another potential way to conduct Bayesian inference for SGM is MCMC methods (Raftery and Lewis, 1996) as SGM has a closed-form solution in the Fourier frequency domain. However, even under an explicit frequency domain solution, the likelihood function of SGM can be complicated (Verma et al., 2022b), which hampers analyzing the properties of the posterior density. Moreover, MCMC methods require a long burn-in step to reach the equilibrium distribution and samples from the equilibrium distribution are correlated. These properties make sampling from MCMC rather time-consuming for SGM. In addition, the computational cost of MCMC methods cannot be amortized which means the time-consuming MCMC procedure needs to be run anew for each observation, regardless of prior observations. A previous MCMC-based inference was unable to capture the spectral features using a nonlinear neural mass model (Xie et al., 2019). In comparison, SBI is more flexible than MCMC methods. Due to the powerful neural network, it can easily handle the complicated likelihood function. More importantly, SBI is trained with simulation samples that help to reduce the requirements of real data. Once the model is trained, it can be applied to new observations without retraining. Therefore, compared with MCMC methods, SBI may be preferable for practical Bayesian inference.

It is worth noting that while this paper focuses on SGM, the SBI approach can be a robust and efficient alternative for parameter estimation of any complex generative model, e.g. above-mentioned coupled neural mass or DCM models. The key trade-off involves whether upfront simulation of a large number of forward model runs is practical and whether there is a compelling use case for achieving rapid inference of an unseen observation.

4.2. Limitations of the current approach

In the current inference procedure, our simulator outputs include the regional PSD and the spatial distribution of alpha band power — together they form a relatively high-dimensional feature space. While Algorithm 1 is capable of handling this, the high dimensionality of output features increases the computational burden and causes difficulty in learning useful information with neural networks from the data. Although we reported some basic diagnostics in Fig. 4 to verify the validity of our inference, the high-dimensional output hampers more extensive posterior diagnostics. Possible workarounds to deal with this issue include extracting some key features from the PSD and spatial distribution manually or embedding a neural network to learn the key features automatically. More experiments are required in this direction. Another limitation is the large number of simulation samples required in SBI-SGM, which slows the inference procedure and increases the computational burden. The number of required simulation samples can be dramatically reduced with multi-round inference (Greenberg et al., 2019) via focusing the training on a particular observation. Although the trained model loses the generality for other observations, it can be very useful when we are only interested in one specific observed dataset.

4.3. Potential applications and future work

In clinical practice, it can sometimes be even more important to know how accurate our estimate is than simply to know the best point estimate (Trkulja and Hraba , 2019). For example, using only point estimates, it can be difficult to compare computational biomarkers from different cohorts. Even if two cohorts have very different values of the

biophysical parameters, no statistically robust conclusion can be drawn without knowing the uncertainty of those estimates. In such cases, SBI-SGM will be extremely helpful as it gives the posterior distributions of the parameters which fully captures the uncertainty of estimates. With posterior distributions, credible intervals and other measures of uncertainty can be easily obtained. This can also be used to obtain population-level parameters that are homogeneous across a population despite the individual variability, which can aid in establishing the descriptive validity of models like SGM (Bassett et al., 2018). Lastly, it can also be used to obtain time-varying posteriors of model parameters that can capture the fast temporal fluctuations in MEG, as has been done previously using point estimates (Verma et al., 2022b).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Data availability

Data will be made available on request.

References

Abdelnour Farras, Dayan Michael, Devinsky Orrin, Thesen Thomas, Raj Ashish, 2018. Functional brain connectivity is predictable from anatomic network's Laplacian eigen-structure. NeuroImage 172, 728–739. 10.1016/j.neuroimage.2018.02.016, URL https://www.sciencedirect.com/science/article/pii/S1053811918301046. [PubMed: 29454104]

Abdelnour Farras, Voss Henning U., Raj Ashish, 2014. Network diffusion accurately models the relationship between structural and functional brain connectivity networks. NeuroImage 90, 335–347. 10.1016/j.neuroimage.2013.12.039, URL https://www.sciencedirect.com/science/article/pii/S1053811913012597. [PubMed: 24384152]

Achard Sophie, Salvador Raymond, Whitcher Brandon, Suckling John, Bullmore Ed, 2006. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical Hubs. J. Neurosci. 26 (1), 63–72. 10.1523/JNEUROSCI.3874-05.2006, arXiv:https://www.jneurosci.org/content/26/1/63.full.pdf URL https://www.jneurosci.org/content/26/1/63. [PubMed: 16399673]

Alonso Leandro M., Marder Eve, 2019. Visualization of currents in neural models with similar behavior and different conductance densities. Elife 8.

Alstott Jeffrey, Breakspear Michael, Hagmann Patric, Cammoun Leila, Sporns Olaf, 2009. Modeling the Impact of Lesions in the Human Brain. PLoS Comput. Biol. 5 (6), e1000408. 10.1371/journal.pcbi.1000408. [PubMed: 19521503]

Anon, 2020. MATLAB version 9.8.0.1451342 (R2020a) Update 5. The Mathworks, Inc., Natick, Massachusetts.

Bassett Danielle Smith, Bullmore Ed, 2006. Small-world brain networks. The Neuroscientist 12 (6), 512–523. 10.1177/1073858406293182,. [PubMed: 17079517]

Bassett Danielle S., Bullmore Edward T., 2009. Human brain networks in health and disease. Curr. Opin. Neurol. 22 (4), 340. [PubMed: 19494774]

Bassett Danielle S., Zurn Perry, Gold Joshua I., 2018. On the nature and use of models in network neuroscience. Nat. Rev. Neurosci. 19 (9), 566–578. [PubMed: 30002509]

Breakspear Michael, 2017. Dynamic models of large-scale brain activity. Nature Neurosci. 20 (3), 340–352. [PubMed: 28230845]

Buckner Randy L., Snyder Abraham Z., Shannon Benjamin J., LaRossa Gina, Sachs Rimmon, Fotenos Anthony F., Sheline Yvette I., Klunk William E., Mathis Chester A., Morris John C., Mintun Mark A., 2005. Molecular, structural, and functional characterization of alzheimer's disease: Evidence for a relationship between default activity, amyloid, and memory. J. Neurosci. 25 (34), 7709–7717. 10.1523/JNEUROSCI.2177-05.2005, arXiv:https://www.jneurosci.org/content/25/34/7709.full.pdf. [PubMed: 16120771]

Bullmore Ed, Sporns Olaf, 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. Nat. Rev. Neurosci. 10 (3), 186–198. [PubMed: 19190637]

Cabral Joana, Kringelbach Morten L., Deco Gustavo, 2014. Exploring the network dynamics underlying brain activity during rest. Progr. Neurobiol. 114, 102–131. 10.1016/j.pneurobio.2013.12.005, URL https://www.sciencedirect.com/science/article/pii/S0301008213001457.

Cabral Joana, Kringelbach Morten L., Deco Gustavo, 2017. Functional connectivity dynamically evolves on multiple time-scales over a static structural connectome: Models and mechanisms. NeuroImage 160, 84–96. 10.1016/j.neuroimage.2017.03.045, URL https://www.sciencedirect.com/science/article/pii/S1053811917302537 Functional Architecture of the Brain. [PubMed: 28343985]

Chatterjee Nivedita, Sinha Sitabhra, 2007. Understanding the mind of a worm: hierarchical network structure underlying nervous system function in c. elegans. In: Banerjee Rahul, Chakrabarti Bikas K. (Eds.), Models of Brain and Mind. In: Progress in Brain Research, vol. 168, Elsevier, pp. 145–153. 10.1016/S0079-6123(07)68012-1, URL https://www.sciencedirect.com/science/article/pii/S0079612307680121.

Cranmer Kyle, Brehmer Johann, Louppe Gilles, 2020. The frontier of simulation-based inference. Proc. Natl. Acad. Sci. 117 (48), 30055–30062. [PubMed: 32471948]

Dalal Sarang S, Zumer JM, Agrawal V, Hild KE, Sekihara Kensuke, Nagarajan SS, 2004. NUTMEG: a neuromagnetic source reconstruction toolbox. Neurol. Clin. Neurophysiol. NCN 2004, 52. [PubMed: 16012626]

David Olivier, Friston Karl J., 2003. A neural mass model for MEG/EEG:: coupling and neuronal dynamics. NeuroImage 20 (3), 1743–1755. 10.1016/j.neuroimage.2003.07.015, URL https://www.sciencedirect.com/science/article/pii/S1053811903004579. [PubMed: 14642484]

Deistler Michael, Goncalves Pedro J., Macke Jakob H., 2022. Truncated proposals for scalable and hassle-free simulation-based inference. arXiv preprint arXiv:2210. 04815.

Desikan Rahul S, Ségonne Florent, Fischl Bruce, Quinn Brian T, Dickerson Bradford C, Blacker Deborah, Buckner Randy L, Dale anders M, Maguire R Paul, Hyman Bradley T, et al. , 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. NeuroImage 31 (3), 968–980. [PubMed: 16530430]

Destexhe Alain, Sejnowski Terrence J., 2009. The Wilson–Cowan model, 36 years later. Biol. Cybernet. 101 (1), 1–2.

Durkan Conor, Bekasov Artur, Murray Iain, Papamakarios George, 2019. Neural spline flows. Adv. Neural Inf. Process. Syst. 32.

Boustani El, Sami Destexhe, Alain, 2009. A master equation formalism for macroscopic modeling of asynchronous irregular activity states. Neural Comput. 21 (1), 46–100. [PubMed: 19210171]

Fan Lingzhong, Li Hai, Zhuo Junjie, Zhang Yu, Wang Jiaojian, Chen Liangfu, Yang Zhengyi, Chu Congying, Xie Sangma, Laird Angela R, et al. , 2016. The human brainnetome atlas: a new brain atlas based on connectional architecture. Cerebral cortex 26 (8), 3508–3526. [PubMed: 27230218]

Fischl Bruce, Salat David H, Busa Evelina, Albert Marilyn, Dieterich Megan, Haselgrove Christian, Kouwe andre Van Der, Killiany Ron, Kennedy David, Klaveness Shuna, Montillo Albert, Makris Nikos, Rosen Bruce, Dale anders M, 2002. Whole Brain Segmentation : Automated Labeling of Neuroanatomical Structures in the Human Brain. Neuron 33, 341–355. [PubMed: 11832223]

Fornito Alex, Zalesky andrew, Breakspear Michael, 2015. The connectomics of brain disorders. Nat. Rev. Neurosci. 16 (3), 159–172. [PubMed: 25697159]

Gabry Jonah, Simpson Daniel, Vehtari Aki, Betancourt Michael, Gelman andrew, 2017. Visualization in Bayesian workflow. arXiv preprint arXiv:1709.01449.

Ghosh A, Rho Y, McIntosh AR, Kötter R, Jirsa VK, 2008. Cortical network dynamics with time delays reveals functional connectivity in the resting brain. Cogn. Neurodyn. 2 (2), 115–120. [PubMed: 19003478]

Gonçalves Pedro J, Lueckmann Jan-Matthis, Deistler Michael, Nonnenmacher Marcel, Öcal Kaan, Bassetto Giacomo, Chintaluri Chaitanya, Podlaski William F, Haddad Sara A, Vogels Tim P, et al. , 2020. Training deep neural density estimators to identify mechanistic models of neural dynamics. Elife 9, e56261. [PubMed: 32940606]

Greenberg David, Nonnenmacher Marcel, Macke Jakob, 2019. Automatic posterior transformation for likelihood-free inference. In: International Conference on Machine Learning. PMLR, pp. 2404–2414.

Haan Willem de, Mott Katherine, Straaten Elisabeth C. W. van, Scheltens Philip, Stam Cornelis J., 2012. Activity Dependent Degeneration Explains Hub Vulnerability in Alzheimer's Disease. PLoS Comput. Biol. 8 (8), e1002582. 10.1371/journal.pcbi.1002582. [PubMed: 22915996]

Hartoyo Agus, Cadusch Peter J., Liley David T.J., Hicks Damien G., 2019. Parameter estimation and identifiability in a neural population model for electro-cortical activity. PLoS Comput. Biol. 15 (5), 1–27. 10.1371/journal.pcbi.1006694.

He Yong, Chen Zhang, Evans Alan, 2008. Structural insights into aberrant topological patterns of large-scale cortical networks in alzheimer's disease. J. Neurosci. 28 (18), 4756–4766. 10.1523/JNEUROSCI.0141-08.2008, URL https://www.jneurosci.org/content/28/18/4756 arXiv:https://www.jneurosci.org/content/28/18/4756.full.pdf. [PubMed: 18448652]

Henderson Harold V., Searle., 1979. Vec and vech operators for matrices, with some uses in Jacobians and multivariate statistics. Canad. J. Statist. 7 (1), 65–81.

Hermundstad Ann M, Bassett Danielle S, Brown Kevin S, Aminoff Elissa M, Clewett David, Freeman Scott, Frithsen Amy, Johnson Arianne, Tipper Christine M, Miller Michael B, et al. , 2013. Structural foundations of resting-state and task-based functional connectivity in the human brain. Proc. Natl. Acad. Sci. 110(15), 6169–6174. [PubMed: 23530246]

Hilbe Joseph M., 2009. Logistic Regression Models. Chapman and hall/CRC.

Honey Christopher J., Sporns Olaf, 2008. Dynamical consequences of lesions in cortical networks. Human Brain Mapping 29 (7), 802–809. 10.1002/hbm.20579. [PubMed: 18438885]

Honey CJ, Sporns O, Cammoun L, Gigandet X, Thiran JP, Meuli R, Hagmann P, 2009. Predicting human resting-state functional connectivity from structural connectivity. Proc. Natl. Acad. Sci. 106 (6), 2035–2040. 10.1073/pnas.0811168106, URL https://www.pnas.org/content/106/6/2035 arXiv:https://www.pnas.org/content/106/6/2035.full.pdf. [PubMed: 19188601]

Jenkinson Mark, Beckmann Christian F., Behrens Timothy E.J., Woolrich Mark W., Smith Stephen M., 2012. FSL. NeuroImage 62 (2), 782–790. [PubMed: 21979382]

Kiebel Stefan J, Garrido Marta I, Moran Rosalyn J, Friston Karl J, 2008. Dynamic causal modelling for EEG and MEG. Cogn. Neurodyn. 2 (2), 121. [PubMed: 19003479]

Leshno Moshe, Lin Vladimir Ya, Pinkus Allan, Schocken Shimon, 1993. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. Neural Netw. 6 (6), 861–867.

Llanas Bernardo, Lantarón Sagrario, Sáinz, Francisco J, 2008. Constructive approximation of discontinuous functions by neural networks. Neural Process. Lett. 27 (3), 209–226.

Lueckmann Jan-Matthis, Boelts Jan, Greenberg David, Goncalves Pedro, Macke Jakob, 2021. Benchmarking simulation-based inference. In: International Conference on Artificial Intelligence and Statistics. PMLR, pp. 343–351.

Marrelec Guillaume, Krainik Alexandre, Duffau Hugues, Pélégrini-Issac Mélanie, Lehéricy Stéphane, Doyon Julien, Benali Habib, 2006. Partial correlation for functional brain interactivity investigation in functional MRI. NeuroImage 32 (1), 228–237. [PubMed: 16777436]

McInnes Leland, Healy John, Melville James, 2018. Umap: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:1802. 03426.

Muldoon Sarah Feldt, Pasqualetti Fabio, Gu Shi, Cieslak Matthew, Grafton Scott T., Vettel Jean M., Bassett Danielle S., 2016. Stimulation-based control of dynamic brain networks. PLoS Comput. Biol. 12 (9), 1–23. 10.1371/journal.pcbi.1005076.

Owen Julia P., Li Yi-Ou, Ziv Etay, Strominger Zoe, Gold Jacquelyn, Bukhpun Polina, Wakahiro Mari, Friedman Eric J., Sherr Elliott H., Mukherjee Pratik, 2013. The structural connectome of the human brain in agenesis of the corpus callosum. NeuroImage 70, 340–355. 10.1016/j.neuroimage.2012.12.031, URL https://www.sciencedirect.com/science/article/pii/S1053811912012165. [PubMed: 23268782]

Papamakarios George, Murray Iain, 2016. Fast ϵ-free inference of simulation models with bayesian conditional density estimation. Adv. Neural Inf. Process. Syst. 29.

Park Hae-Jeong, Friston Karl, 2013. Structural and functional brain networks: from connections to cognition. Science 342 (6158).

Pereira Inês, Frässle Stefan, Heinzle Jakob, Schöbi Dario, Do Cao Tri, Gruber Moritz, Stephan Klaas E., 2021. Conductance-based dynamic causal modeling: A mathematical review of its application to cross-power spectral densities. NeuroImage 245, 118662. 10.1016/j.neuroimage.2021.118662, URL https://www.sciencedirect.com/science/article/pii/S1053811921009356. [PubMed: 34687862]

Persson K, Björkman M, 1988. Annoyance due to low frequency noise and the use of the d'B (A) scale. J. Sound Vib. 127 (3), 491–497.

Pinotsis Dimitris A., Moran Rosalyn J., Friston Karl J., 2012. Dynamic causal modeling with neural fields. NeuroImage 59 (2), 1261–1274. [PubMed: 21924363]

Raftery Adrian E., Lewis Steven M., 1996. Implementing mcmc. Markov chain Monte Carlo practice 115–130.

Raj Ashish, Cai Chang, Xie Xihe, Palacios Eva, Owen Julia, Mukherjee Pratik, Nagarajan Srikantan, 2020. Spectral graph theory of brain oscillations. Hum. Brain Map. 41 (11), 2980–2998, arXiv:10.1002/hbm.24991.

Raj Ashish, Verma Parul, Nagarajan Srikantan, 2022. Structure-function models of temporal, spatial, and spectral characteristics of non-invasive whole brain functional imaging.. Front. Neurosci. 16, 959557. [PubMed: 36110093]

Ritter Petra, Schirner Michael, McIntosh Anthony R., Jirsa Viktor K., 2013. The virtual brain integrates computational modeling and multimodal neuroimaging. Brain Connect. 3 (2), 121–145. 10.1089/brain.2012.0120, [PubMed: 23442172]

Rubinov Mikail, Sporns Olaf, van Leeuwen Cees, Breakspear Michael, 2009. Symbiotic relationship between brain structure and dynamics. BMC Neurosci. 10 (1), 1–18. [PubMed: 19126204]

Sanz-Leon Paula, Knock Stuart A., Spiegler andreas, Jirsa Viktor K., 2015. Mathematical framework for large-scale brain network modeling in The Virtual Brain. NeuroImage 111, 385–430. 10.1016/j.neuroimage.2015.01.002, URL https://www.sciencedirect.com/science/article/pii/S1053811915000051. [PubMed: 25592995]

Sengupta Biswa, Friston Karl J., Penny Will D., 2015. Gradient-free MCMC methods for dynamic causal modelling. NeuroImage 112, 375–381. [PubMed: 25776212]

Sengupta Biswa, Friston Karl J., Penny Will D., 2016. Gradient-based MCMC samplers for dynamic causal modelling. NeuroImage 125, 1107–1118. [PubMed: 26213349]

Siettos Constantinos, Starke Jens, 2016. Multiscale modeling of brain dynamics: from single neurons and networks to mathematical tools. WIREs Syst. Biol. Med. 8 (5), 438–458. 10.1002/wsbm.1348.

Singh Matthew F., Braver Todd S., Cole Michael W., Ching ShiNung, 2020. Estimation and validation of individualized dynamic brain models with resting state fMRI. NeuroImage

221, 117046. 10.1016/j.neuroimage.2020.117046, URL https://www.sciencedirect.com/science/article/pii/S1053811920305322. [PubMed: 32603858]

Spiegler A, Jirsa V, 2013. Systematic approximations of neural fields through networks of neural masses in the virtual brain. NeuroImage 83, 704–725. 10.1016/j.neuroimage.2013.06.018, URL https://www.sciencedirect.com/science/article/pii/S1053811913006551. [PubMed: 23774395]

Strogatz Steven H., 2001. Exploring complex networks. Nature 410 (6825), 268–276. [PubMed: 11258382]

Suárez Laura E., Markello Ross D., Betzel Richard F., Misic Bratislav, 2020. Linking structure and function in macroscale brain networks. Trends in Cognitive Sciences 24 (4), 302–315. 10.1016/j.tics.2020.01.008, URL https://www.sciencedirect.com/science/article/pii/S1364661320300267. [PubMed: 32160567]

Tejero-Cantero Alvaro, Boelts Jan, Deistler Michael, Lueckmann Jan-Matthis, Durkan Conor, Gonçalves Pedro J., Greenberg David S., Macke Jakob H., 2020. Sbi: A toolkit for simulation-based inference. J. Open Sour. Softw. 5 (52), 2505. 10.21105/joss.02505.

Tomm Christian, Avermann Michael, Vogels Tim, Gerstner Wulfram, Petersen Carl, 2011. The influence of structure on the response properties of biologically plausible neural network models. BMC Neurosci. 12 (1), 1–2. [PubMed: 21208416]

Trkulja Vladimir, Hraba Pero, 2019. Confidence intervals: what are they to us, medical doctors? Croatian Med. J. 60 (4), 375.

van den Heuvel Martijn P., Mandl René C.W., Kahn René S., Hulshoff Pol Hilleke E., 2009. Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. Hum. Brain Map. 30 (10), 3127–3141. 10.1002/hbm.20737, URL 10.1002/hbm.20737 arXiv:10.1002/hbm.20737.

Geit Van, Werner Gevaert, Michael Chindemi, Giuseppe Rössert, Christian Courcol, Jean-Denis Muller, Eilif B Schürmann, Felix Segev, Idan Markram, Henry, 2016. BluePyOpt: leveraging open source software and cloud infrastructure to optimise model parameters in neuroscience. Front. Neuroinform. 10, 17. [PubMed: 27375471]

Verma Parul, Nagarajan Srikantan, Raj Ashish, 2022a. Spectral graph theory of brain oscillations —-Revisited and improved. NeuroImage 249, 118919. 10.1016/j.neuroimage.2022.118919, URL https://www.sciencedirect.com/science/article/pii/S1053811922000490. [PubMed: 35051584]

Verma Parul, Nagarajan Srikantan, Raj Ashish, 2022b. Stability and dynamics of a spectral graph model of brain oscillations. Netw. Neurosci. 1–43. 10.1162/netn_a_00263, arXiv:10.1162/netn_a_00263/2034454/netn_a_00263.pdf. [PubMed: 35350585]

Wilson Hugh R., Cowan Jack D., 1973. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. Kybernetik 13 (2), 55–80. [PubMed: 4767470]

Xie Xihe, Kuceyeski Amy, Shah Sudhin A, Schiff Nicholas D, Nagarajan SS, Raj Ashish, 2019. Parameter identifiability and non-uniqueness in connectome based neural mass models. 480012, bioRxiv.

Xie Xi He, Stanley Megan, Damasceno Pablo F., 2020. Raj-Lab-UCSF/spectrome: Spectral Graph Model of Neural Dynamics on Connectomes. 10.5281/zenodo.3620935.

Yang Genevieve J., Murray John D., Wang Xiao-Jing, Glahn David C., Pearlson Godfrey D., Repovs Grega, Krystal John H., Anticevic Alan, 2016. Functional hierarchy underlies preferential connectivity disturbances in schizophrenia. Proc. Natl. Acad. Sci. 113 (2), E219–E228. 10.1073/pnas.1508436113. [PubMed: 26699491]

Zimmermann J, Perry A, Breakspear M, Schirner M, Sachdev P, Wen W, Kochan NA, Mapstone M, Ritter P, McIntosh AR, Solodkin A, 2018. Differentiation of Alzheimer's disease based on local and global parameters in personalized Virtual Brain models. NeuroImage Clin. 19, 240–251. [PubMed: 30035018]
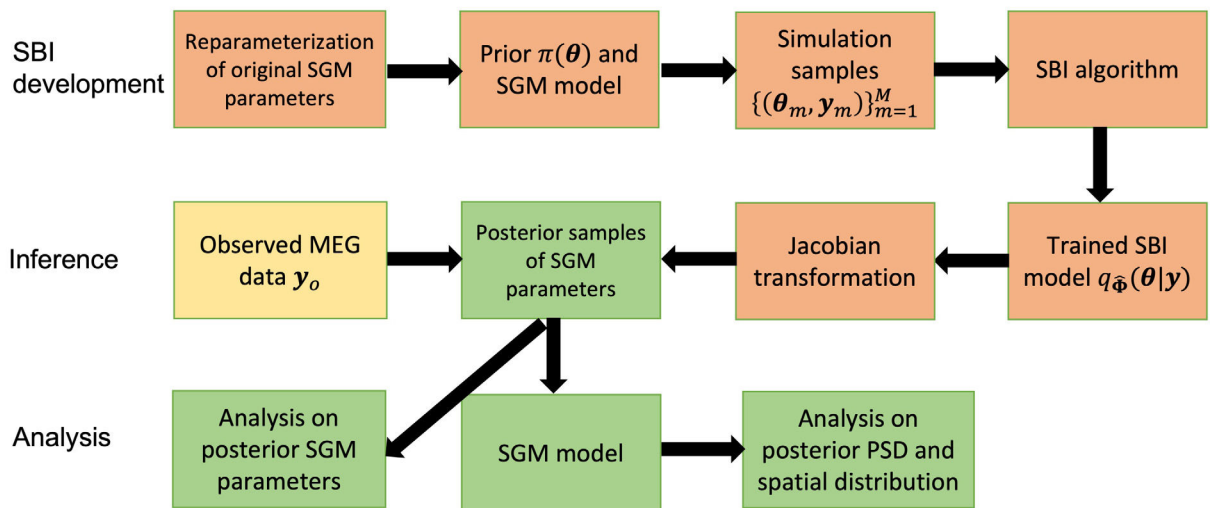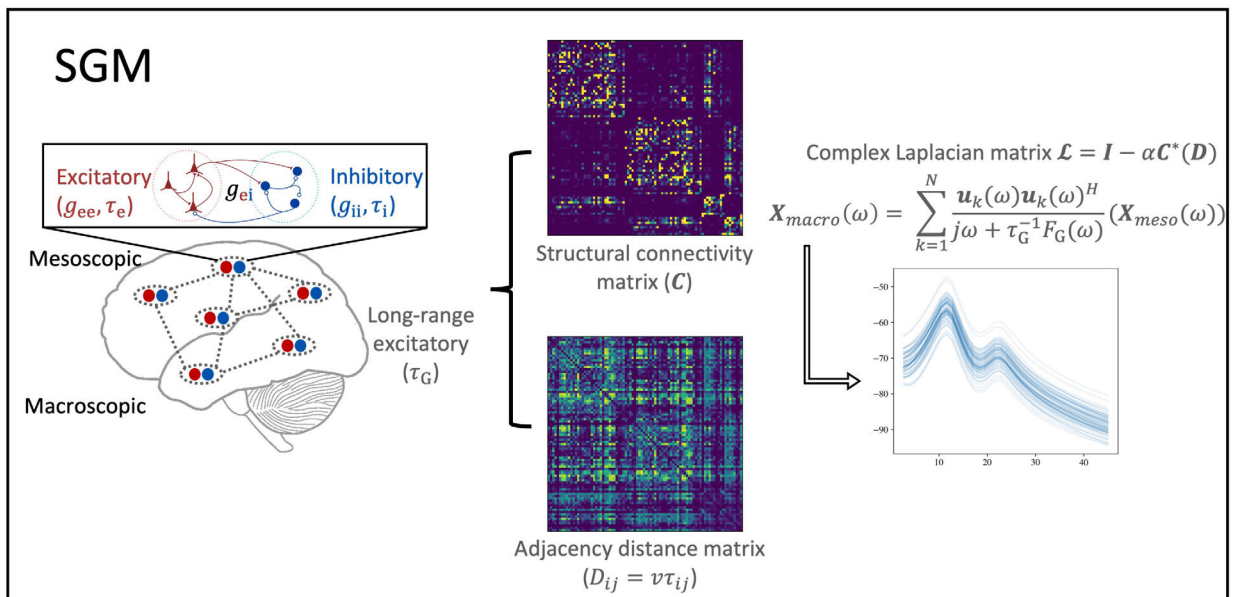
Fig. 1.
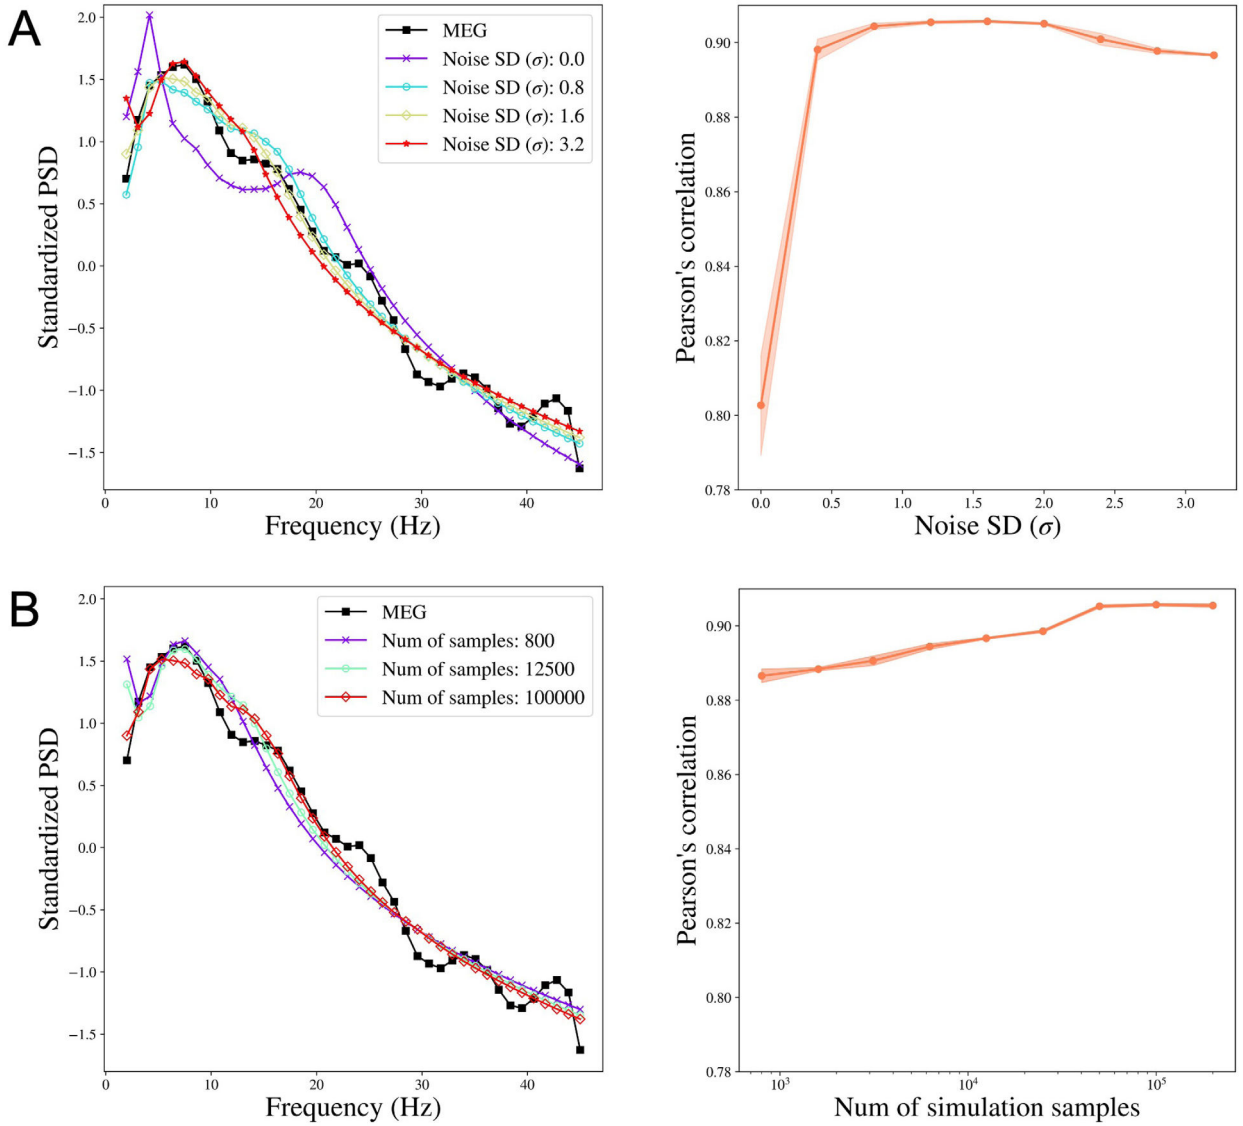Top: A schematic of SGM. Bottom: The pipeline of SBI inference for SGM.

Fig. 2.
The performance of SBI-SGM when varying the noise SD and the number of simulation samples. **A:** Left: Median standardized PSD obtained from MEG and SBI-SGM with different noise SDs. Right: Change of Pearson's correlation between reconstructed average PSD and the observed PSD when varying noise SDs. The red shadow indicates its 95% confidence interval. **B:** Left: Median standardized PSD obtained from MEG and SBI-SGM with different number of simulation samples. Right: Change of Pearson's correlation between reconstructed average PSD and the observed PSD when varying the number of simulation samples. The red shadow indicates its 95% confidence interval.
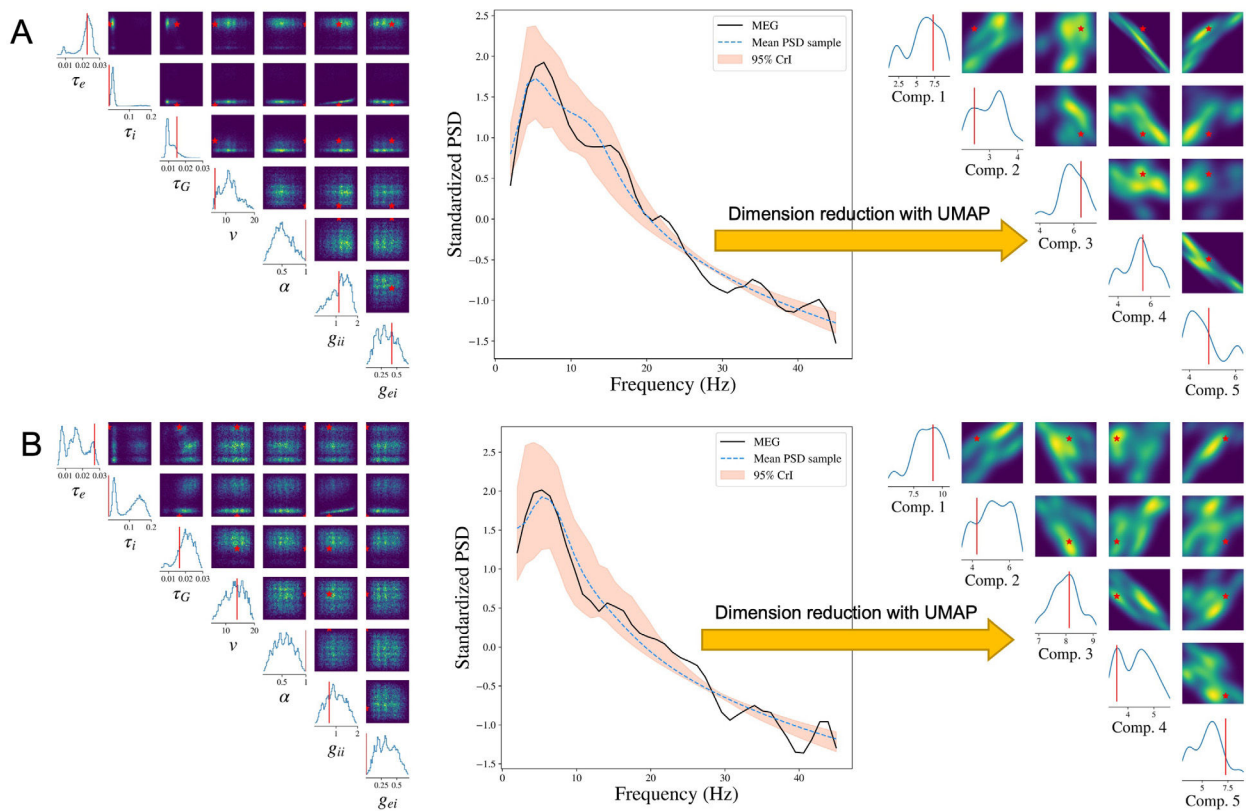
Fig. 3.
Result analysis from SBI-SGM for two representative MEG data whose Pearson's correlations between reconstructed PSD and the observed PSD are the top two closest to the median correlation across 36 subjects. **A/B:** Left: Posterior density of 7 parameters for one subject in the MEG dataset. The red vertical line and red star indicate the location of the point estimate from the annealing algorithm. Middle: Posterior mean PSD and the 95% credible interval for the subject. The black curve indicates the observed average PSD across ROIs. Right: Density estimations and observed values of low-dimensional representations after mapping raw vectorized PSDs to a 5-dimensional embedding manifold with UMAP. The red vertical line and red star indicate the location of the representation for observed MEG data in the manifold.
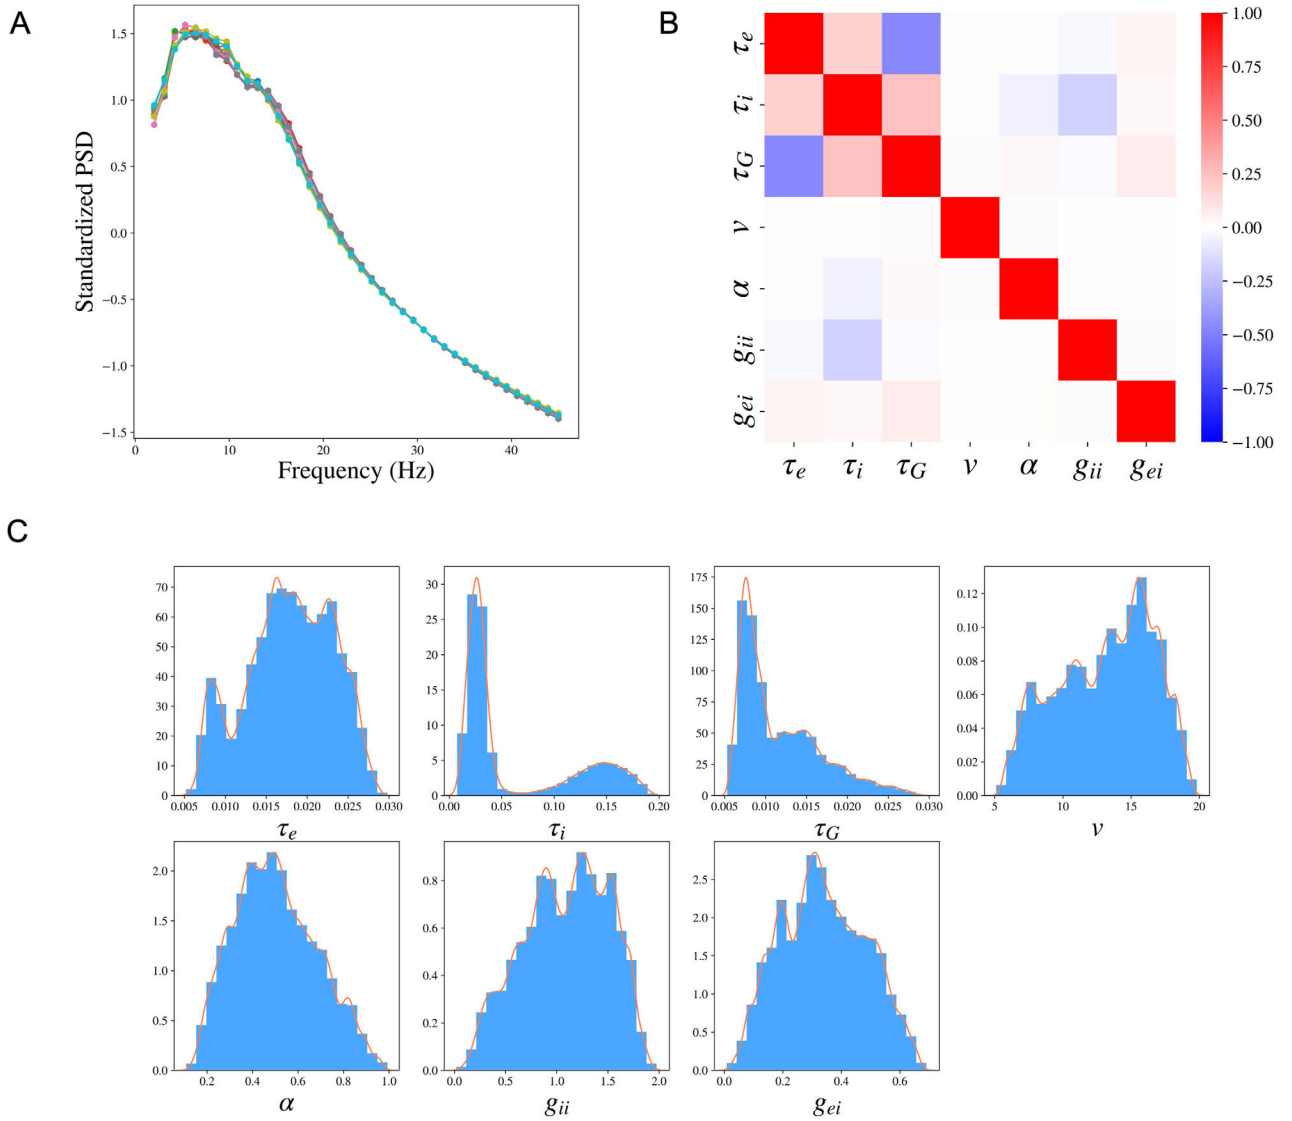
Fig. 4.
Cohort level analysis with SBI-SGM across the 36 MEG datasets. **A:** The median standardized PSDs obtained from SBI-SGM in 10 repetitions under noise SD 1.6 and number of simulation samples 100,000. The correlations between reconstructed and observed PSDs are between 0.905 and 0.907 in the 10 repetitions. **B:** Partial correlation between each pair of parameters averaged over 36 subjects. **C:** Histograms and the corresponding kernel density estimations of the posterior SGM parameters.
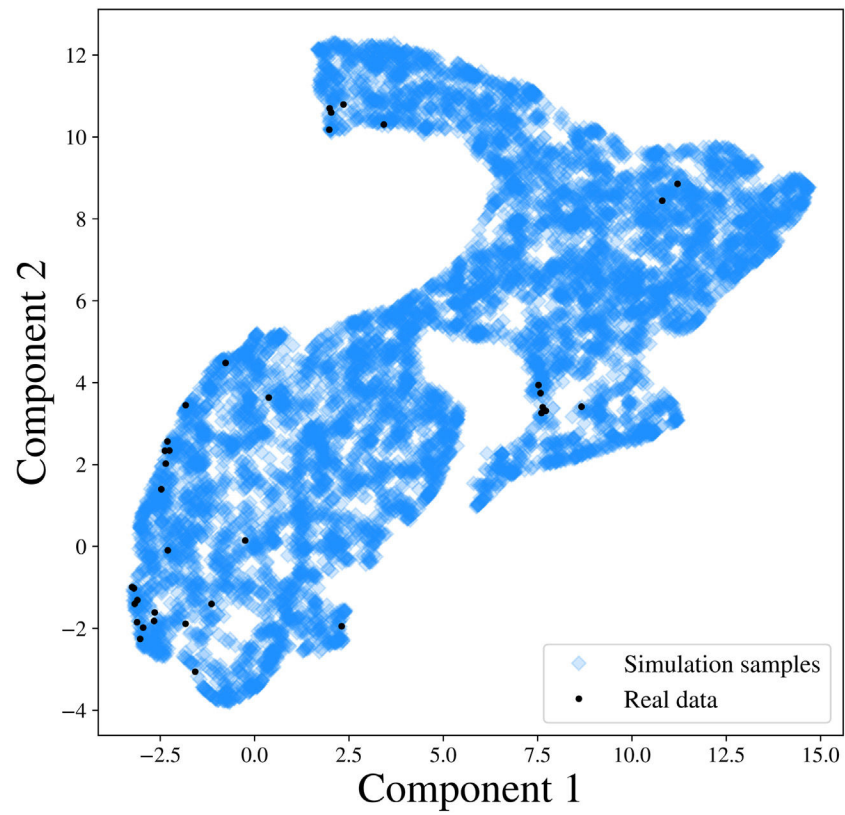
Fig. 5.
PSD representations after vectorizing and mapping the simulated PSDs and 36 observed PSDs to a 2-dimensional embedding manifold with the UMAP. In this instance, we adopt UMAP to reduce the dimensionality of the vectorized PSD matrix to 2 for visualization purposes.
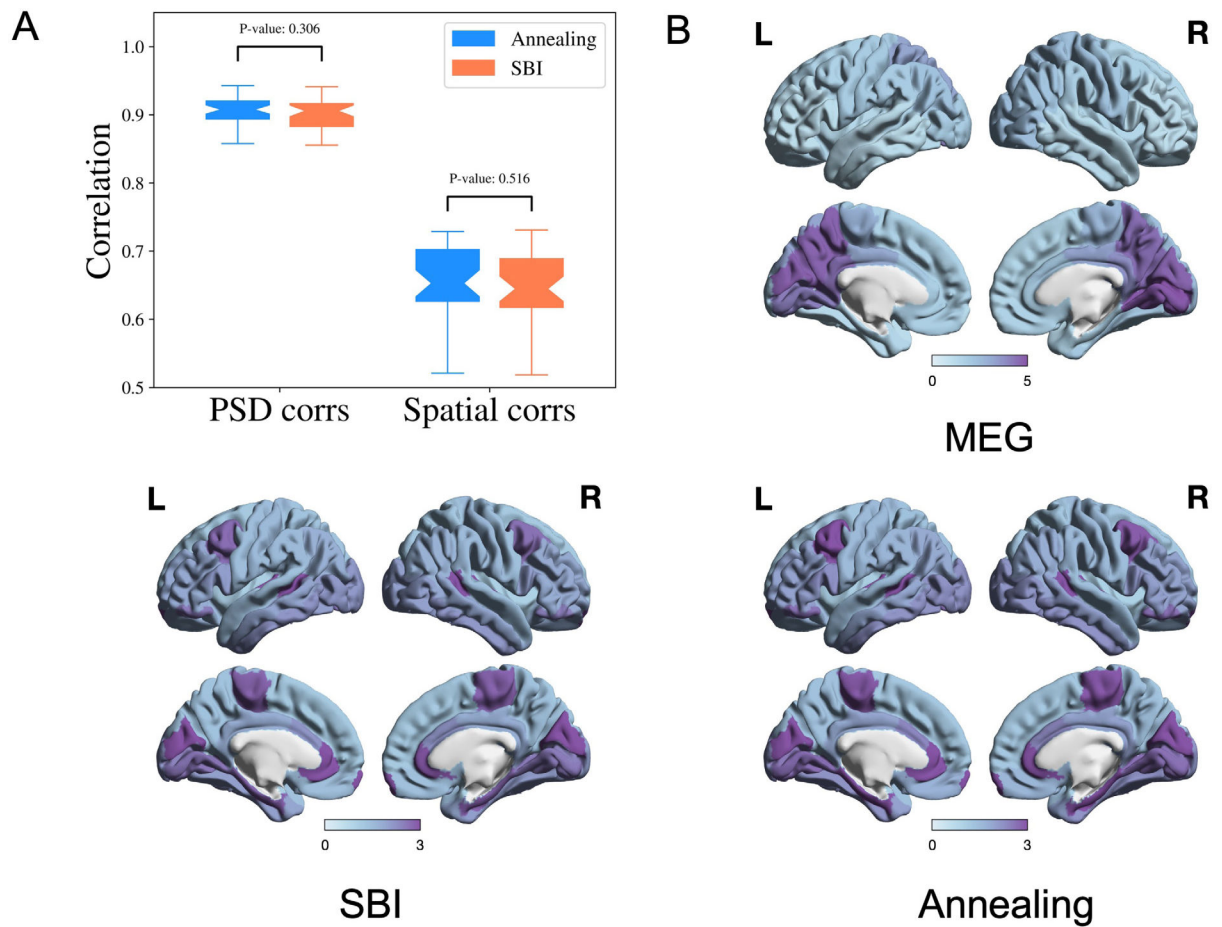
Fig. 6.
Performance of SBI and annealing on SGM is comparable. **A:** Pearson's correlation of PSD from each ROI and spatial correlation for the alpha frequency band. P-values are from two-sample t-tests. **B:** Comparison of the observed and reconstructed spatial distributions from the SBI and annealing algorithms of the alpha frequency band, averaged over all the subjects.

Table 1

SGM parameters and bounds for the parameter estimation for SBI and annealing.

Name	Symbol	Lower/upper bound
Excitatory time constant	τ_e	[0.005 s, 0.03 s]
Inhibitory time constant	τ_i	[0.005 s, 0.2 s]
Long-range connectivity coupling constant	α	[0.1, 1]
Transmission speed	v	[5 m/s, 20 m/s]
Alternating population gain	g_{ei}	[0.001, 0.7], [0.001, 0.5], [0.001, 0.4]
Inhibitory gain	g_{ii}	[0.001, 2.0], [0.001, 1.5], [0.001, 1.5]
Graph time constant	τ_G	[0.005 s, 0.03 s]