Bayesian Inference of a Spectral Graph Model for Brain Oscillations

Huaqing Jin[1, *], Parul Verma[1, *], Fei Jiang[2], Srikantan Nagarajan[1], Ashish Raj[1]

[1]Department of Radiology and Biomedical Imaging
University of California San Francisco, USA San Francisco, CA
[2] Department of Epidemiology and Biostatistics
University of California San Francisco, USA San Francisco, CA

[*]co-first

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 wide-band 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 can not 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 [1, 2]. Functional activity between the grey 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 (DTI) from MRI. The brain structure-function (SC-FC) relationship is then investigated using various data-driven and mathematical modeling-based techniques, assuming structural connectivity (SC) 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 [3–17] as well as modeling techniques [18–28] have been employed to uncover the brain SC-FC relationships, 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 [29–34]. For example, Zimmermann et al. [33] 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 confidence bounds 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 [35]. 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.

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 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 were highlighted in previous studies [36–38], 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 wide-band power spectral density (PSD) as well as the spatial distribution of the alpha band power obtained from MEG [39]. 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 [40, 41]. 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 [36–38].

In our prior works we had estimated SGM parameters using a global optimization algorithm, the dual annealing method, as point estimates [39–41, 37], as a preferred alternative over continuous gradient descent-based minimization. However, since the objective function of estimating SGM parameter [39, 40] 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 [42]. 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. Finally, it allows us to benefit from the availability of an unbounded number of simulations, thereby helping overcome the critical issue of lack of large sets of empirical data in medical settings.

Using the SBI tool applied to the SGM model, we demonstrate that the model 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 SC-FC mapping, 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 Magnetoencephalography (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 [39, 43]. Data collection procedure was described in [39, 40]. 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. MEG recordings were collected for 5 minutes while the subjects were resting and had their eyes closed. Out of the 5-minute recording, a 1-minute snippet was chosen which was most noise free. 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. MEG recordings were downsampled to 600 Hz, followed by a bandpass filtering of the signals between 2 to 45 Hz using `firls` in MATLAB [44] and generation of the static frequency PSD for every region of interest using the `pmtm` algorithm in MATLAB [44].

2.2 Spectral graph model

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 [39–41]. 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 [39, 40].

SGM is characterized by eight parameters, which include the excitatory and inhibitory time constants τ_e, τ_i and neural gains g_{ee}, 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. The neural gain g_{ee} is kept as 1 to ensure parameter identifiability [39], so the parameters of interest to be inferred in SGM are $\mathbf{s} = (\tau_e, \tau_i, \alpha, v, g_{ei}, g_{ii}, \tau_G)^{\mathrm{T}}$. Given the signals with N regions of interest (ROIs) in the time domain is $[x_1(t), \ldots, x_N(t)]^{\mathrm{T}}$, the closed-form solution of SGM is obtained in the Fourier domain:

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

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, $\mathcal{F}$ is the Fourier transformation, $\mathcal{L}$ 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 [39].

2.3 Simulation-based inference for SGM

Simulation-based inference (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 [45–47]. We adapt the SBI method for SGM parameter estimation and inference (referred to as SBI-SGM). Let $\mathbf{X}(\mathbf{s},\Omega) = \{\mathbf{X}(\mathbf{s},\omega)\}_{\omega\in\Omega}$ be the model output PSD in dB scale[48] 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

$$\mathbf{y} = G\left\{\mathbf{X}(\mathbf{s},\Omega)\right\} + \boldsymbol{\epsilon}, \tag{2}$$

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

The random noise in (2) captures the biological noise, artifacts, and measurement error in the MEG data. Without this random noise, the target posterior density is discontinuous, which is difficult to estimate due to the well-known Gibbs phenomenon [49]. 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 [50].

Similarly, we define the observed samples $\mathbf{y}_o = G\left\{\mathbf{X}_{\text{MEG}}(\Omega)\right\}$ 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 [51]. To address this issue, we re-parameterize the parameters so that the posterior sampling can be performed on the real line. More specifically, let $\mathcal{H}(x)$ be a scaled logit transformation function [52] defined as

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

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

Under a Bayesian framework, we are interested in the posterior distribution of $\boldsymbol{\theta}$ given $\mathbf{y}$, particularly the Bayesian credible interval of $\boldsymbol{\theta}$, which captures the uncertainty of the SGM parameters. To obtain the credible interval, we estimate the posterior distribution of $\boldsymbol{\theta}$ through the SBI procedure [42]. The density of $\mathbf{y}$ is denoted by $p(\mathbf{y}|\boldsymbol{\theta})$ following (2), which is a multivariate Gaussian density function. We impose a multivariate Gaussian prior $\pi(\boldsymbol{\theta})$ on $\boldsymbol{\theta}$. The posterior density is $q_{\boldsymbol{\Phi}}(\boldsymbol{\theta}|\mathbf{y}) \propto \pi(\boldsymbol{\theta})p(\mathbf{y}|\boldsymbol{\theta})$, where $\boldsymbol{\Phi}$ is the unknown parameters that determine the posterior distribution. Instead of obtaining the posterior density for the SGM parameters $\mathbf{s}$ directly, we first derive the posterior density for $\boldsymbol{\theta}$, which results in the target posterior distribution through a Jacobian transformation [53].

We use a deep learning architecture, namely neural spline flow (NSF) [47], to model the functional form of $q_{\boldsymbol{\Phi}}$, where $\boldsymbol{\Phi}$ is the parameters in the deep learning network. The dimension of $\boldsymbol{\Phi}$ increases with the number of network layers in NSF, and when the dimension of $\boldsymbol{\Phi}$ approaches infinity, $q_{\boldsymbol{\Phi}}$ approaches the true posterior distribution. When the deep learning architecture is given, $\boldsymbol{\Phi}$ is the only unknown parameter in $q_{\boldsymbol{\Phi}}(\boldsymbol{\theta}|\mathbf{y})$. Hence estimating the posterior density is equivalent to estimating $\boldsymbol{\Phi}$. Now note that the true posterior distribution maximizes $\mathbf{E}[\log\{q_{\boldsymbol{\Phi}}(\boldsymbol{\theta}|\mathbf{y})\}]$, where the expectation is taken with respect to $\mathbf{y}$ and $\boldsymbol{\theta}$, we propose to obtain an estimator for $\boldsymbol{\Phi}$ through maximizing the empirical version of $\mathbf{E}[\log\{q_{\boldsymbol{\Phi}}(\boldsymbol{\theta}|\mathbf{y})\}]$, that is $\frac{1}{M}\sum_{m=1}^{M}\log\left\{q_{\boldsymbol{\Phi}}(\boldsymbol{\theta}_m|\mathbf{y}_m)\right\}$, where the samples $\mathbf{y}_m$ and $\boldsymbol{\theta}_m$, $m = 1, \ldots, M$ are the simulated realizations of $\mathbf{y}$ and $\boldsymbol{\theta}$ based on $p(\mathbf{y}|\boldsymbol{\theta})$ and $\pi(\boldsymbol{\theta})$, respectively.

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

2.4 Implementation details

For each subject, we use their MEG data from 68 cortical regions according to the Desikan–Killiany atlas [54] to obtain the posterior samples of the SGM parameters $\mathbf{s}$. Under this atlas, we obtain a 68-region $\mathbf{X}(\mathbf{s}, \omega)$ at frequency ω, and the dimension of $G\left\{\mathbf{X}(\mathbf{s}, \Omega)\right\}$ is 2788.

We implement SBI-SGM using the sbi package in Python (https://www.mackelab.org/sbi/) [55], 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 Simulation 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

Algorithm 1 Posterior estimation with re-parameterization

```
Require: A multivariate Gaussian prior $\pi(\boldsymbol{\theta}) \sim N(0, 100\mathbf{I})$ a likelihood $p(\mathbf{y}|\boldsymbol{\theta})$, an observation $\mathbf{y}_o$.
  Simulation:
  for $m = 1, \ldots, M$ do
      Sample $\boldsymbol{\theta}_m \sim \pi(\boldsymbol{\theta})$
      Sample $\mathbf{y}_m \sim p(\mathbf{y}|\boldsymbol{\theta}_m)$
  end for
  Optimization:
  $\widehat{\boldsymbol{\Phi}} \leftarrow \underset{\boldsymbol{\Phi}}{\operatorname{argmin}} -\frac{1}{M}\sum_{m=1}^{M} \log\{q_{\boldsymbol{\Phi}}(\boldsymbol{\theta}_m|\mathbf{y}_m)\}$
  return $q_{\widehat{\boldsymbol{\Phi}}}(\boldsymbol{\theta}|\mathbf{y}_o)$ as the estimate of the posterior distribution of $\boldsymbol{\theta}$. The posterior distribution of s is $q_{\widehat{\boldsymbol{\Phi}}} \times |\det(\mathbf{J})|$, where $\mathbf{J}$ is the Jacobian matrix, i.e., $\mathbf{J} = \partial\boldsymbol{\theta}/\partial\mathrm{s}$.
```

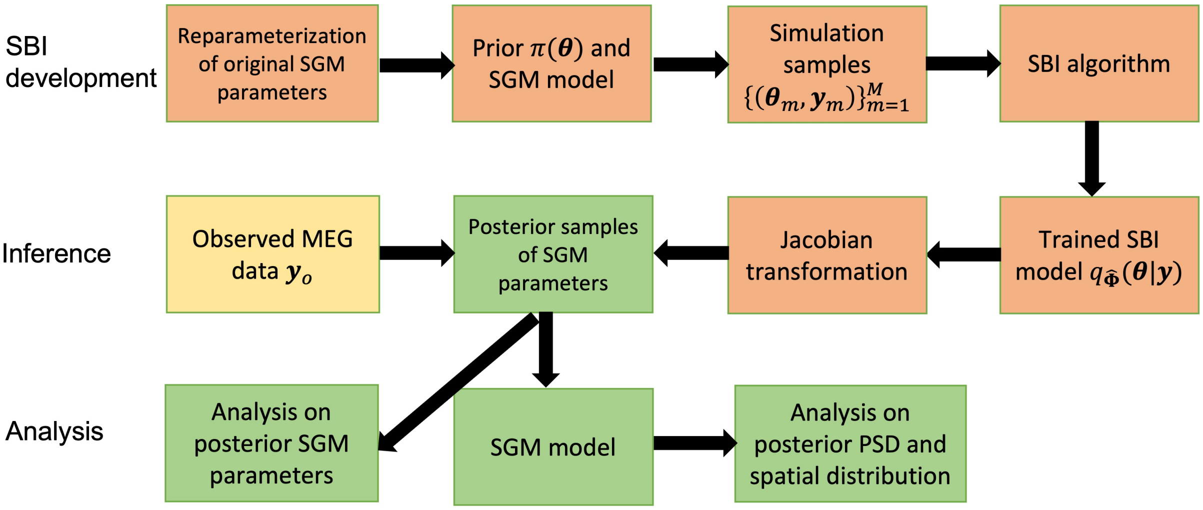

Figure 1: The pipeline of SBI inference for SGM.

density for each observed $\mathbf{y}_o$, we draw a posterior sample of SGM parameters, denoted by $\widetilde{\mathrm{s}}$. We then obtain $\mathbf{X}(\widetilde{\mathrm{s}}, \Omega)$ using (1). Finally, we construct the standardized PSD and spatial distribution of the alpha band PSD as a G transformation of $\mathbf{X}(\widetilde{\mathrm{s}}, \Omega)$, 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 Figure 1.

We compare the performance of SBI-SGM with the performance of the annealing SGM approach [39–41], namely Ann-SGM, on our MEG data. The details of the annealing implementation can be found in [40]. 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 [41] 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 [41]. For the other five parameters, SBI-SGM uses the same bounds as Ann-SGM does.

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]
Excitatory gain	g_{ee}	n/a

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 Simulation 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.

Figure 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 Simulation step in Algorithm 1. Figure 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 Figure 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 Simulation step in Algorithm 1 in the subsequent analyses.

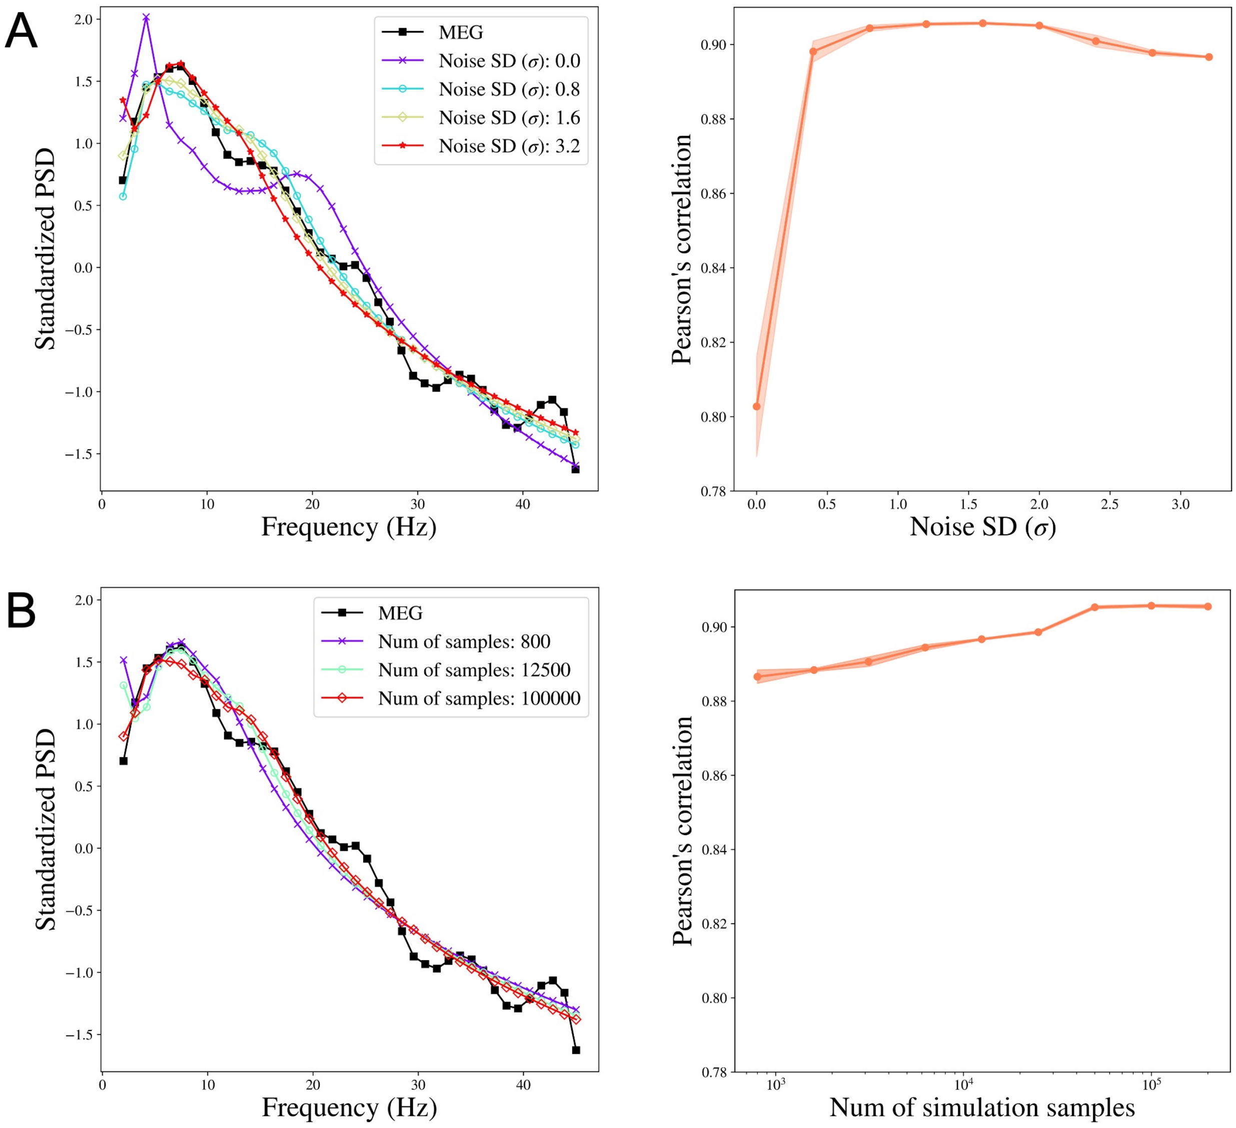

Figure 2: The performance of SBI-SGM when varying the noise SD and the number of simulation samples. **A:** Left: Median standardized power spectral density (PSD) obtained from MEG and SBI-SGM with different noise SDs. Right: Change of Pearson's correlation between reconstructed average PSD and the observed 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.

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

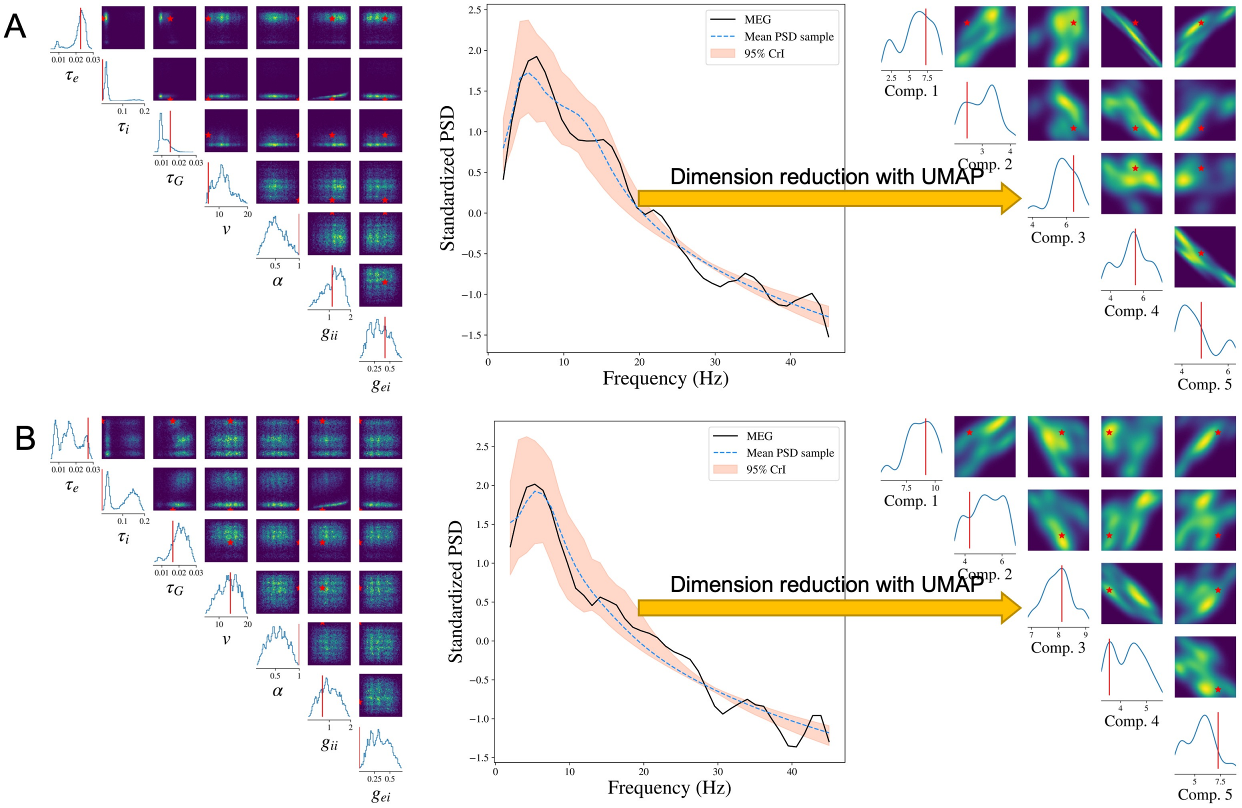

Figure 3: Result analysis from SBI-SGM for two representative MEG data. **A/B:** Left: Posterior density of 7 parameters for one subject in the MEG dataset. The red vertical line and red star indicate the location of the point estimate from the annealing algorithm. Middle: Posterior mean PSD and the 95% credible interval for the subject. The black curve indicates the observed average PSD. Right: Density estimations and observed values of low-dimensional representations after mapping raw PSDs to a 5-dimensional embedding manifold with the uniform manifold approximation and projection (UMAP). The red vertical line and red star indicate the location of the representation for observed MEG data in the manifold.

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 Figures 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 Figure 3 shows the posterior mean and 95% credible interval (CrI) of the PSD from SBI-SGM. In each subject, the 95% CrI 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 [40].

Moreover, we project the reconstructed and observed PSDs and map them onto a 5-dimensional manifold using the uniform manifold approximation and projection (UMAP) method proposed by [56]. As shown in the right panels of Figure 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 [57].

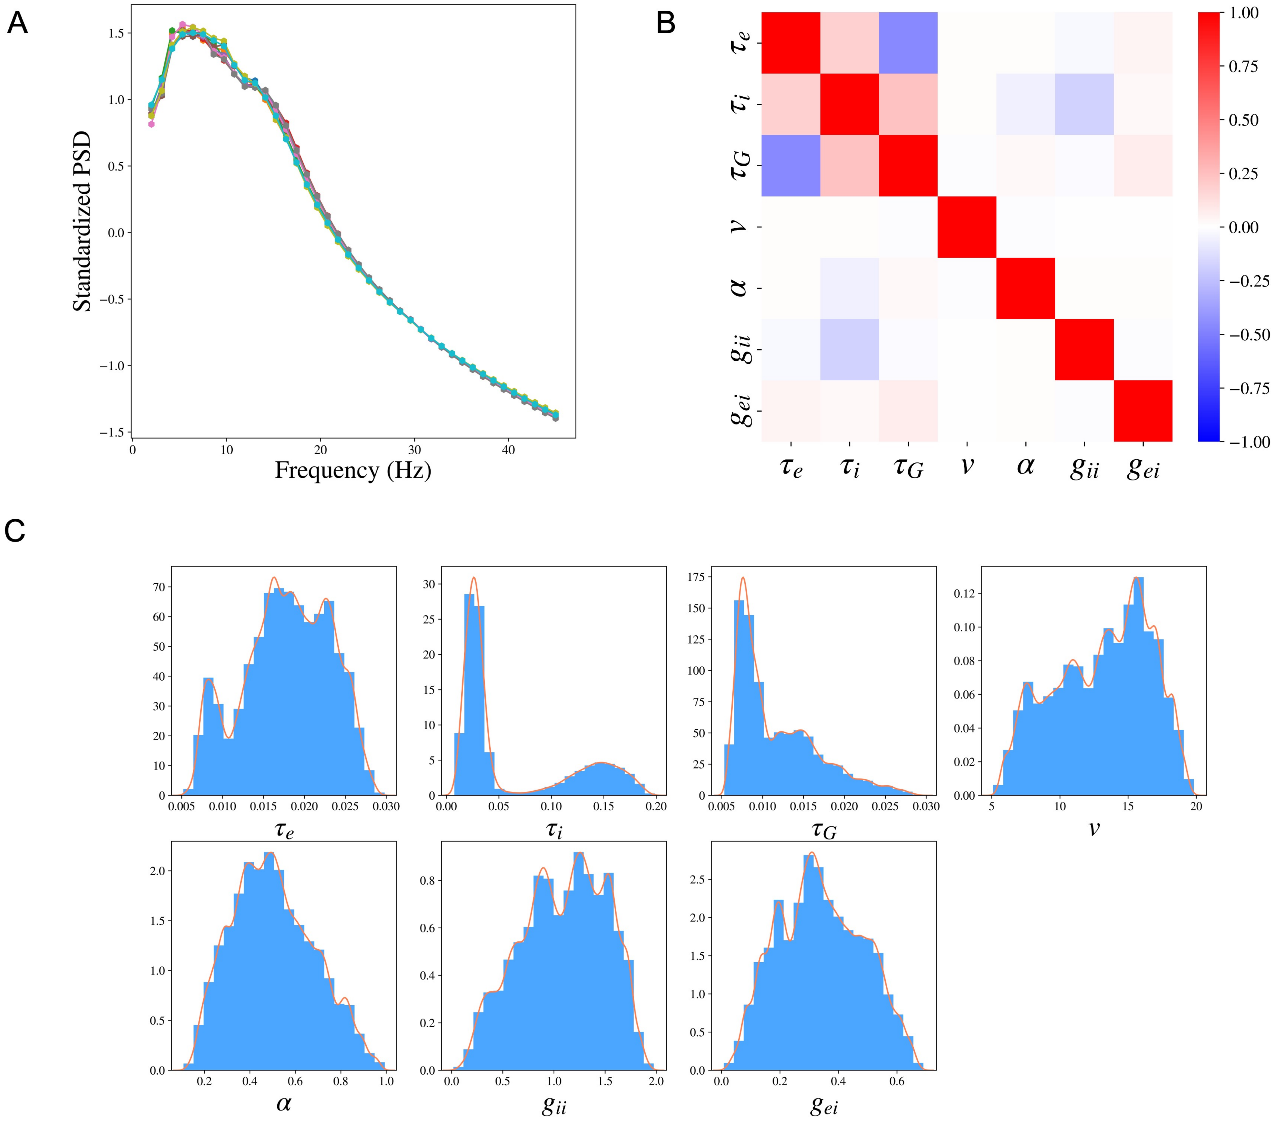

Figure 4: Cohort level SBI-SGM across the 36 MEG datasets. **A:** The median standardized power spectral densities (PSDs) obtained from SBI-SGM in 10 repetitions under noise SD 1.6 and number of simulation samples $100,000$. The correlations between reconstructed and observed PSDs are between 0.905 and 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.

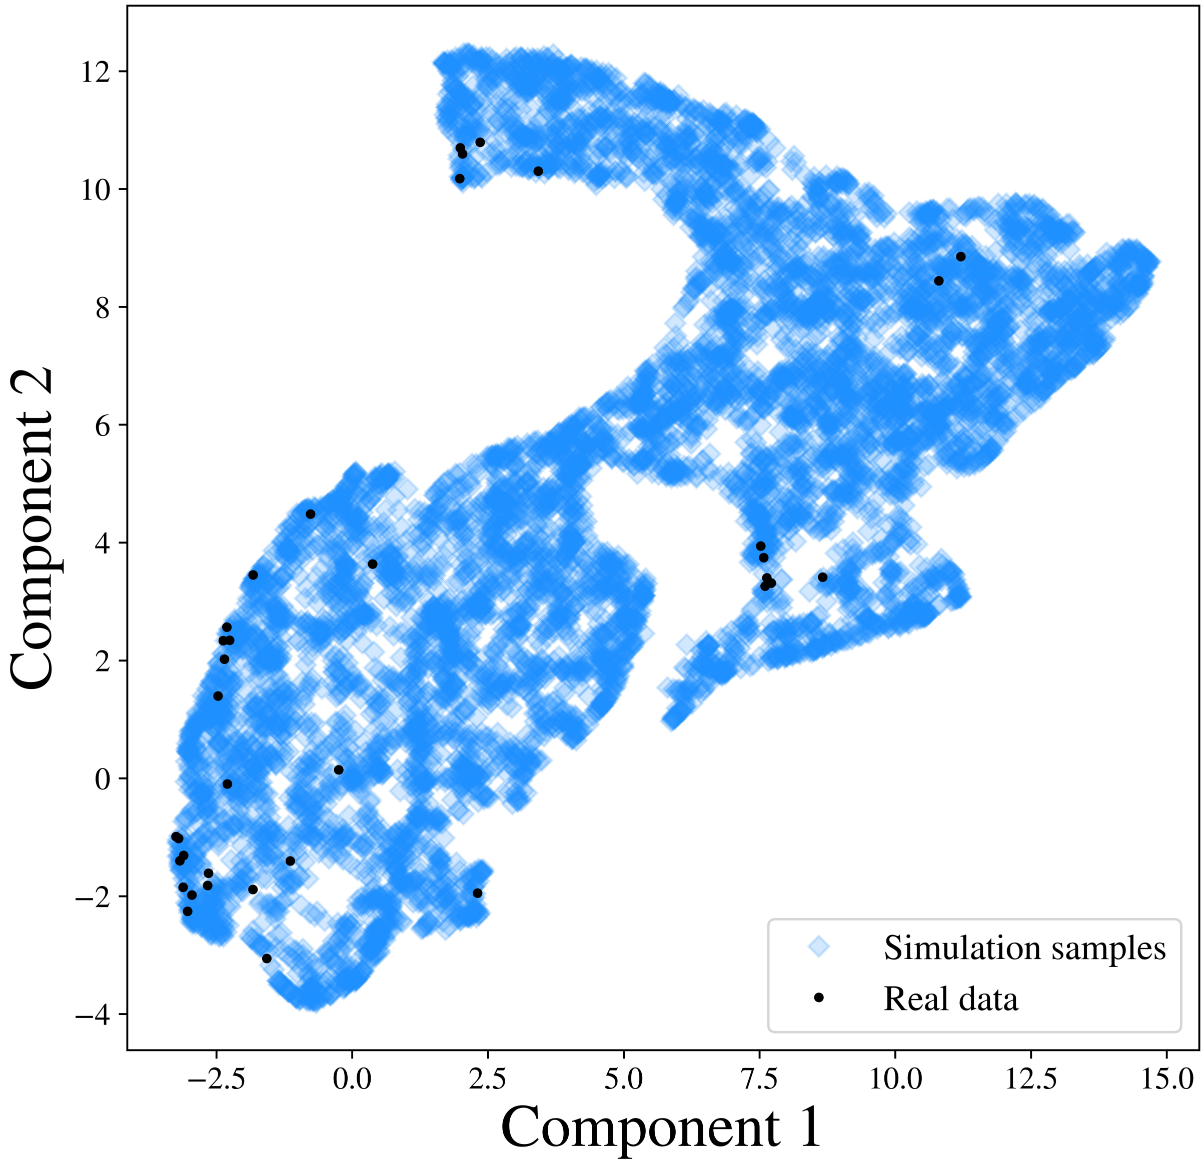

Figure 5: PSD representations after mapping the simulated PSDs and 36 observed PSDs to a 2-dimensional embedding manifold with the UMAP.

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 Figure 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 [58], which examines the correlation between any given two parameters after removing the effect from other parameters, Figure 4B shows the pair-wise partial correlation averaged over 36 subjects. As shown in Figure 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. Figure 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 $\mathbf{s}$, and obtain simulated PSDs through (1). We then compare the simulated PSDs with the observed ones. To facilitate the visualization, we utilize the UMAP method to project the simulated samples of PSDs and observed PSDs to a 2-dimensional embedding manifold. As shown in Figure 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 likelihood of the data.

3.5 SBI-SGM and Ann-SGM comparision

We compare SBI-SGM with Ann-SGM. In Figure 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 [37].

As shown in Figure 6A, SBI-SGM gives similar average correlation and spatial correlation as Ann-SGM does with insignificant p-values from two-sample t-tests. In Figure 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.

One notable advantage of SBI-SGM is that it is much faster than Ann-SGM. When a template connectome is given, SBI-SGM provides a universal posterior distribution that can be applied to all the 36 MEG data after one training. SBI-SGM takes approximately 2 hours to accomplish the Bayesian inference on SGM parameters for all subjects. On the other hand, Ann-SGM takes about 8 hours for each subject and the algorithm needs to repeat for each subject. Parallel computing can further improve the computational speed of SBI-SGM. In conclusion, SBI-SGM has a similar performance as Ann-SGM on recovering observed PSD and spatial distribution from the alpha band but is much more computationally efficient.

4 Discussion

Models with complex and stochastic simulators have been extensively applied in many areas of science and engineering [59]. In neuroscience, such computational models are typically built via incorporating biological

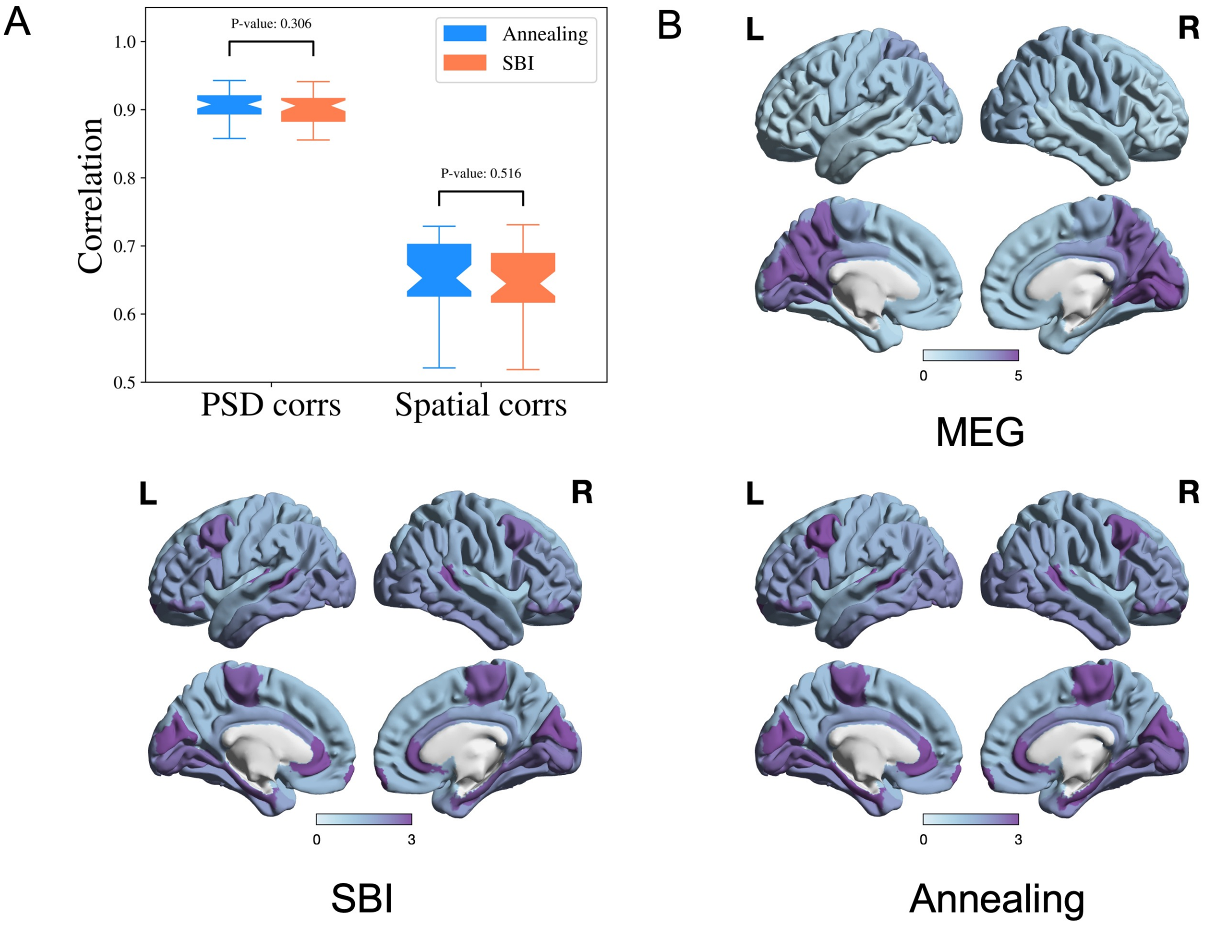

Figure 6: Performance of SBI and annealing on SGM is comparable. **A:** Pearson's correlation of PSD from each ROI and spatial correlation for the alpha frequency band. P-values are from two-sample t-tests. **B:** Comparison of the observed and reconstructed spatial distributions from the SBI and annealing algorithms of the alpha frequency band, averaged over all the subjects

mechanisms and hypothetical intuitions to explain the observed phenomena inferred from the neuroimaging data [35]. 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 [60], genetic algorithm [61] or simulated annealing [39, 40].

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 [35, 62]. 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 [42, 47]. For this reason, SBI may be challenging for coupled non-linear models such as NMMs and the Virtual Brain [63] 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 [24, 64, 27, 26], 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 [47, 42], 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 [65, 66]. 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 [67]. 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 [37], and parameter identifiability is not guaranteed [36]. These challenges have been discussed in detail elsewhere [38, 37]. 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 global model parameters.

Another potential way to conduct Bayesian inference for SGM is Markov chain Monte Carlo (MCMC) methods [68] 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 [41], 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 can not 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 [38]. 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 Figure 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 [46] 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 [69]. 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 [70]. 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 [41].

Data availability

The code and processed datasets for this work can be found in this github repository: https://github.com/JINhuaqing/SBI-SGM.

Acknowledgments

This work was supported by NIH grants R01NS092802, R01NS183412, R01AG062196, R01AG072753, R01EB022717, R01DC013979, R01NS100440, R01DC176960, R01DC017091, K25AG071840, UCOP-MRP-17-454755, Alzheimer's Association grant AARFD-22-923931, and an industry research contract from Ricoh MEG Inc. The template Human Connectome Project (HCP) connectome used in the preparation of this work was obtained from the MGH-USC HCP database (https://ida.loni.usc.edu/login.jsp). The HCP project's MGH-USC Consortium (Principal Investigators: Bruce R. Rosen, Arthur W. Toga and Van Wedeen; U01MH093765) is supported by the NIH Blueprint Initiative for Neuroscience Research Grant; the National Institutes of Health grant P41EB015896; and the Instrumentation Grants S10RR023043, 1S10RR023401, 1S10RR019307. Collectively, the HCP is the result of efforts of co-investigators from the University of Southern California, Martinos Center at Massachusetts General Hospital (MGH), Washington University, and the University of Minnesota.

References

[1] Alex Fornito, Andrew Zalesky, and Michael Breakspear. The connectomics of brain disorders. *Nature Reviews Neuroscience*, 16(3):159–172, 2015.

[2] Laura E. Suárez, Ross D. Markello, Richard F. Betzel, and Bratislav Misic. Linking structure and function in macroscale brain networks. *Trends in Cognitive Sciences*, 24(4):302–315, 2020.

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

[4] Randy L. Buckner, Abraham Z. Snyder, Benjamin J. Shannon, Gina LaRossa, Rimmon Sachs, Anthony F. Fotenos, Yvette I. Sheline, William E. Klunk, Chester A. Mathis, John C. Morris, and Mark A. Mintun. Molecular, structural, and functional characterization of alzheimer's disease: Evidence for a relationship between default activity, amyloid, and memory. *Journal of Neuroscience*, 25(34):7709–7717, 2005.

[5] Sophie Achard, Raymond Salvador, Brandon Whitcher, John Suckling, and Ed Bullmore. A Resilient, Low-Frequency, Small-World Human Brain Functional Network with Highly Connected Association Cortical Hubs. *Journal of Neuroscience*, 26(1):63–72, 2006.

[6] Danielle Smith Bassett and Ed Bullmore. Small-world brain networks. *The Neuroscientist*, 12(6):512–523, 2006. PMID: 17079517.

[7] Nivedita Chatterjee and Sitabhra Sinha. Understanding the mind of a worm: hierarchical network structure underlying nervous system function in c. elegans. In Rahul Banerjee and Bikas K. Chakrabarti,

editors, *Models of Brain and Mind*, volume 168 of *Progress in Brain Research*, pages 145–153. Elsevier, 2007.

[8] Yong He, Zhang Chen, and Alan Evans. Structural insights into aberrant topological patterns of large-scale cortical networks in alzheimer's disease. *Journal of Neuroscience*, 28(18):4756–4766, 2008.

[9] A Ghosh, Y Rho, AR McIntosh, R Kötter, and VK Jirsa. Cortical network dynamics with time delays reveals functional connectivity in the resting brain. *Cognitive Neurodynamics*, 2(2):115–120, 2008.

[10] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198, 2009.

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

[12] Martijn P. van den Heuvel, René C.W. Mandl, René S. Kahn, and Hilleke E. Hulshoff Pol. Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Human Brain Mapping*, 30(10):3127–3141, 2009.

[13] Danielle S Bassett and Edward T Bullmore. Human brain networks in health and disease. *Current Opinion in Neurology*, 22(4):340, 2009.

[14] Farras Abdelnour, Henning U. Voss, and Ashish Raj. Network diffusion accurately models the relationship between structural and functional brain connectivity networks. *NeuroImage*, 90:335–347, 2014.

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

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

[17] Farras Abdelnour, Michael Dayan, Orrin Devinsky, Thomas Thesen, and Ashish Raj. Functional brain connectivity is predictable from anatomic network's Laplacian eigen-structure. *NeuroImage*, 172:728–739, 2018.

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

[19] Olivier David and Karl J. Friston. A neural mass model for MEG/EEG:: coupling and neuronal dynamics. *NeuroImage*, 20(3):1743–1755, 2003.

[20] Alain Destexhe and Terrence J Sejnowski. The Wilson–Cowan model, 36 years later. *Biological Cybernetics*, 101(1):1–2, 2009.

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

[22] C. J. Honey, O. Sporns, L. Cammoun, X. Gigandet, J. P. Thiran, R. Meuli, and P. Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009.

[23] A. Spiegler and V. Jirsa. Systematic approximations of neural fields through networks of neural masses in the virtual brain. *NeuroImage*, 83:704–725, 2013.

[24] Joana Cabral, Morten L. Kringelbach, and Gustavo Deco. Exploring the network dynamics underlying brain activity during rest. *Progress in Neurobiology*, 114:102–131, 2014.

[25] Sarah Feldt Muldoon, Fabio Pasqualetti, Shi Gu, Matthew Cieslak, Scott T. Grafton, Jean M. Vettel, and Danielle S. Bassett. Stimulation-Based Control of Dynamic Brain Networks. *PLOS Computational Biology*, 12(9):1–23, 09 2016.

[26] Constantinos Siettos and Jens Starke. Multiscale modeling of brain dynamics: from single neurons and networks to mathematical tools. *WIREs Systems Biology and Medicine*, 8(5):438–458, 2016.

[27] Joana Cabral, Morten L. Kringelbach, and Gustavo Deco. Functional connectivity dynamically evolves on multiple time-scales over a static structural connectome: Models and mechanisms. *NeuroImage*, 160:84–96, 2017. Functional Architecture of the Brain.

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

[29] Christopher J. Honey and Olaf Sporns. Dynamical consequences of lesions in cortical networks. *Human Brain Mapping*, 29(7):802–809, 2008.

[30] Jeffrey Alstott, Michael Breakspear, Patric Hagmann, Leila Cammoun, and Olaf Sporns. Modeling the Impact of Lesions in the Human Brain. *PLOS Computational Biology*, 5(6):e1000408, June 2009.

[31] Willem de Haan, Katherine Mott, Elisabeth C. W. van Straaten, Philip Scheltens, and Cornelis J. Stam. Activity Dependent Degeneration Explains Hub Vulnerability in Alzheimer's Disease. *PLOS Computational Biology*, 8(8):e1002582, August 2012.

[32] Genevieve J. Yang, John D. Murray, Xiao-Jing Wang, David C. Glahn, Godfrey D. Pearlson, Grega Repovs, John H. Krystal, and Alan Anticevic. Functional hierarchy underlies preferential connectivity disturbances in schizophrenia. *Proceedings of the National Academy of Sciences*, 113(2):E219–E228, January 2016.

[33] J. Zimmermann, A. Perry, M. Breakspear, M. Schirner, P. Sachdev, W. Wen, N. A. Kochan, M. Mapstone, P. Ritter, A. R. McIntosh, and A. Solodkin. Differentiation of Alzheimer's disease based on local and global parameters in personalized Virtual Brain models. *NeuroImage: Clinical*, 19:240–251, 2018.

[34] Matthew F. Singh, Todd S. Braver, Michael W. Cole, and ShiNung Ching. Estimation and validation of individualized dynamic brain models with resting state fMRI. *NeuroImage*, 221:117046, November 2020.

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

[36] Agus Hartoyo, Peter J. Cadusch, David T. J. Liley, and Damien G. Hicks. Parameter estimation and identifiability in a neural population model for electro-cortical activity. *PLOS Computational Biology*, 15(5):1–27, 05 2019.

[37] Ashish Raj, Parul Verma, and Srikantan Nagarajan. Structure-function models of temporal, spatial, and spectral characteristics of non-invasive whole brain functional imaging. *Frontiers in Neuroscience*, 16:959557–959557, 2022.

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

[39] Ashish Raj, Chang Cai, Xihe Xie, Eva Palacios, Julia Owen, Pratik Mukherjee, and Srikantan Nagarajan. Spectral graph theory of brain oscillations. *Human Brain Mapping*, 41(11):2980–2998, 2020.

[40] Parul Verma, Srikantan Nagarajan, and Ashish Raj. Spectral graph theory of brain oscillations—-revisited and improved. *NeuroImage*, 249:118919, 2022.

[41] Parul Verma, Srikantan Nagarajan, and Ashish Raj. Stability and dynamics of a spectral graph model of brain oscillations. *Network Neuroscience*, pages 1–43, 07 2022.

[42] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.

[43] Xi He Xie, Megan Stanley, and Pablo F. Damasceno. Raj-Lab-UCSF/spectrome: Spectral Graph Model of Neural Dynamics on Connectomes, January 2020.

[44] The Mathworks, Inc., Natick, Massachusetts. *MATLAB version 9.8.0.1451342 (R2020a) Update 5*, 2020.

[45] George Papamakarios and Iain Murray. Fast ε-free inference of simulation models with bayesian conditional density estimation. *Advances in Neural Information Processing Systems*, 29, 2016.

[46] David Greenberg, Marcel Nonnenmacher, and Jakob Macke. Automatic posterior transformation for likelihood-free inference. In *International Conference on Machine Learning*, pages 2404–2414. PMLR, 2019.

[47] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. *Advances in Neural Information Processing Systems*, 32, 2019.

[48] K Persson and M Björkman. Annoyance due to low frequency noise and the use of the db (a) scale. *Journal of Sound and Vibration*, 127(3):491–497, 1988.

[49] Bernardo Llanas, Sagrario Lantarón, and Francisco J Sáinz. Constructive approximation of discontinuous functions by neural networks. *Neural Processing Letters*, 27(3):209–226, 2008.

[50] Moshe Leshno, Vladimir Ya Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks*, 6(6):861–867, 1993.

[51] Michael Deistler, Pedro J Goncalves, and Jakob H Macke. Truncated proposals for scalable and hassle-free simulation-based inference. *arXiv preprint arXiv:2210.04815*, 2022.

[52] Joseph M Hilbe. *Logistic regression models*. Chapman and hall/CRC, 2009.

[53] Harold V Henderson and SR Searle. Vec and vech operators for matrices, with some uses in jacobians and multivariate statistics. *Canadian Journal of Statistics*, 7(1):65–81, 1979.

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

[55] Alvaro Tejero-Cantero, Jan Boelts, Michael Deistler, Jan-Matthis Lueckmann, Conor Durkan, Pedro J. Gonccalves, David S. Greenberg, and Jakob H. Macke. sbi: A toolkit for simulation-based inference. *Journal of Open Source Software*, 5(52):2505, 2020.

[56] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

[57] Jonah Gabry, Daniel Simpson, Aki Vehtari, Michael Betancourt, and Andrew Gelman. Visualization in bayesian workflow. *arXiv preprint arXiv:1709.01449*, 2017.

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

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

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

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

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

[63] Petra Ritter, Michael Schirner, Anthony R. McIntosh, and Viktor K. Jirsa. The virtual brain integrates computational modeling and multimodal neuroimaging. *Brain Connectivity*, 3(2):121–145, 2013. PMID: 23442172.

[64] Paula Sanz-Leon, Stuart A. Knock, Andreas Spiegler, and Viktor K. Jirsa. Mathematical framework for large-scale brain network modeling in The Virtual Brain. *NeuroImage*, 111:385–430, 2015.

[65] Stefan J Kiebel, Marta I Garrido, Rosalyn J Moran, and Karl J Friston. Dynamic causal modelling for eeg and meg. *Cognitive Neurodynamics*, 2(2):121, 2008.

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

[67] Inês Pereira, Stefan Frässle, Jakob Heinzle, Dario Schöbi, Cao Tri Do, Moritz Gruber, and Klaas E. Stephan. Conductance-based dynamic causal modeling: A mathematical review of its application to cross-power spectral densities. *NeuroImage*, 245:118662, 2021.

[68] Adrian E Raftery and Steven M Lewis. Implementing mcmc. *Markov chain Monte Carlo in practice*, pages 115–130, 1996.

[69] Vladimir Trkulja and Pero Hrabavc. Confidence intervals: what are they to us, medical doctors? *Croatian Medical Journal*, 60(4):375, 2019.

[70] Danielle S Bassett, Perry Zurn, and Joshua I Gold. On the nature and use of models in network neuroscience. *Nature Reviews Neuroscience*, 19(9):566–578, 2018.

Supplementary

Spectral graph model

Notation All the vectors and matrices are written in boldface and the scalars are written in normal font. The frequency f of a signal is specified in Hertz (Hz), and the corresponding angular frequency $\omega = 2\pi f$ is used to obtain the Fourier transforms. The connectivity matrix is defined as $\mathbf{C} = c_{jk}$, where c_{jk} is the connectivity strength between regions j and k, normalized by the row degree.

Mesoscopic model

Given region k out of N regions, we denote the local excitatory signal as $x_e(t)$, local inhibitory signal as $x_i(t)$, and the long-range macroscopic signals as $x_k(t)$. Combining the decay of individual signals, coupling of excitatory and inhibitory signals as well as input white Gaussian noise, the evolution models of $x_e(t)$ and $x_i(t)$ are:

$$\frac{\mathrm{d}x_e(t)}{\mathrm{d}t} = -\frac{f_e(t)}{\tau_e} \star \big(g_{ee}\, x_e(t) - g_{ei}\, f_i(t) \star x_i(t)\big) + p(t)\,, \text{ and,} \tag{3}$$

$$\frac{\mathrm{d}x_i(t)}{\mathrm{d}t} = -\frac{f_i(t)}{\tau_i} \star \big(g_{ii}\, x_i(t) + g_{ei}\, f_e(t) \star x_e(t)\big) + p(t)\,, \tag{4}$$

where $f_e(t)$ and $f_i(t)$ are the ensemble average neural impulse response function, $\star$ stands for convolution, $p(t)$ is input noise, parameters g_{ee}, g_{ii}, g_{ei} are neural gain terms, and parameters τ_e, τ_i are characteristic time constants, which are shared for every region k. We assume Gamma-shaped $f_e(t)$ and $f_i(t)$ as

$$f_e(t) = \frac{1}{\tau_e^2} \exp\left(\frac{-t}{\tau_e}\right) \text{ and } f_i(t) = \frac{1}{\tau_i^2} \exp\left(\frac{-t}{\tau_i}\right).$$

Macroscopic model

Accounting for long-range connections between brain regions, the macroscopic signal x_k is assumed to conform to the following evolution model:

$$\frac{\mathrm{d}x_k(t)}{\mathrm{d}t} = -\frac{1}{\tau_G} f_G(t) \star x_k(t) + \frac{\alpha}{\tau_G} f_G(t) \star \sum_{j=1}^{N} c_{jk} x_j(t-\tau^v_{jk}) + \big(x_e(t)+x_i(t)\big)\,, \tag{5}$$

where, τ_G is the graph characteristic time constant, α is the global coupling constant, c_{jk} are elements of the connectivity matrix, τ^v_{jk} is the delay in signals reaching from the j^{th} to the k^{th} region, v is the cortico-cortical fiber conduction speed with which the signals are transmitted. The delay τ^v_{jk} is calculated as d_{jk}/v, where d_{jk} is the distance between regions j and k and $x_e(t)+x_i(t)$ is the input signal determined from Equations (3) and (4).

SGM only includes 8 global parameters as listed in Table 1. The neural gain g_{ee} is kept as 1 to ensure parameter identifiability. Thus, there are only 7 parameters required to be estimated to determine SGM.

Closed-form model solution in the Fourier domain

A salient feature of SGM is that it provides a closed-form solution of brain oscillations under the frequency domain. Let $\mathcal{F}$ be the Fourier transform at angular frequency $\omega = 2\pi f$. Note that the mesoscopic models for different regions share the same parameters, therefore, without loss of generality, we can drop the subscript k.

The solutions for $x_e(t)$ and $x_i(t)$ under the frequency domain are

$$X_e(\omega) = \mathcal{F}(x_e(t)) = \frac{\left\{1 + \frac{g_{ei}F_e(\omega)F_i(\omega)/\tau_e}{j\omega + g_{ii}F_i(\omega)/\tau_i}\right\} P(\omega)}{j\omega + g_{ee}F_e(\omega)/\tau_e + \frac{(g_{ei}F_e(\omega)F_i(\omega))^2}{\tau_e\tau_i(j\omega + g_{ii}F_i(\omega)/\tau_i)}} = H_e(\omega)P(\omega),$$

and

$$X_i(\omega) = \mathcal{F}(x_i(t)) = \frac{\left\{1 + \frac{g_{ei}F_e(\omega)F_i(\omega)/\tau_i}{j\omega + g_{ee}F_e(\omega)/\tau_e}\right\} P(\omega)}{j\omega + g_{ii}F_i(\omega)/\tau_i + \frac{(g_{ei}F_e(\omega)F_i(\omega))^2}{\tau_e\tau_i(j\omega + g_{ee}F_e(\omega)/\tau_e)}} = H_i(\omega)P(\omega),$$

where $P(\omega)$, $F_e(\omega)$, $F_i(\omega)$ are the Fourier transform of $p(t)$, $f_e(t)$ and $f_i(t)$ at angular frequency ω.

We define the complex Laplacian matrix $\mathcal{L}(\omega) = \mathbf{I} - \alpha\mathbf{C}^*(\omega)$ where $\mathbf{C}^*(\omega) = [c_{ij}\exp(-j\omega\tau^v_{ij})]_{i,j=1,\ldots,N}$. The solution of the macroscopic signals at a angular frequency ω is

$$\mathbf{X}(\omega) = [\,\mathcal{F}(x_1(t)), \cdots, \mathcal{F}(x_N(t))]^{\mathrm{T}} = \left(j\omega + \frac{1}{\tau_{\mathbf{G}}} F_{\mathbf{G}}(\omega)\mathcal{L}(\omega)\right)^{-1} H_{\mathrm{local}}(\omega)\mathbf{P}(\omega), \tag{6}$$

where $H_{\mathrm{local}}(\omega) = H_e(\omega) + H_i(\omega)$.

As SGM provides a closed-form solution $\mathbf{X}(\omega)$, we can compare the modeled and empirical power spectra to estimate the global parameters.