Virtual Brain Inference (VBI): A flexible and integrative toolkit for efficient probabilistic inference on virtual brain models

Abolfazl Ziaeemehr[a,*], Marmaduke Woodman[a], Lia Domide[b], Spase Petkoski[a], Viktor Jirsa [†a,*], Meysam Hashemi [†a,*]

[a] *Aix Marseille Univ, INSERM, INS, Inst Neurosci Syst, Marseille, France*
[b] *Codemart, Cluj-Napoca, Romania*

Abstract

Network neuroscience has proven essential for understanding the principles and mechanisms underlying complex brain (dys)function and cognition. In this context, whole-brain network modeling–also known as virtual brain modeling–combines computational models of brain dynamics (placed at each network node) with individual brain imaging data (to coordinate and connect the nodes), advancing our understanding of the complex dynamics of the brain and its neurobiological underpinnings. However, there remains a critical need for automated model inversion tools to estimate control (bifurcation) parameters at large scales associated with neuroimaging modalities, given their varying spatio-temporal resolutions. This study aims to address this gap by introducing a flexible and integrative toolkit for efficient Bayesian inference on virtual brain models, called Virtual Brain Inference (`VBI`). This open-source toolkit provides fast simulations, taxonomy of feature extraction, efficient data storage and loading, and probabilistic machine learning algorithms, enabling biophysically interpretable inference from non-invasive and invasive recordings. Through in-silico testing, we demonstrate the accuracy and reliability of inference for commonly used whole-brain network models and their associated neuroimaging data. `VBI` shows potential to improve hypothesis evaluation in network neuroscience through uncertainty quantification, and contribute to advances in precision medicine by enhancing the predictive power of virtual brain models.

Keywords: Virtual brain models, Simulation-based inference, Structural and functional neuroimaging modalities, Low-dimensional data features

Abbreviations: VBI, virtual brain inference; BOLD, blood-oxygen-level-dependent; fMRI, functional magnetic resonance imaging; EEG, Electroencephalography; MEG, Magnetoencephalography; sEEG, Stereoelectroencephalography; SC, structural connectivity; FC, functional connectivity; FCD, functional connectivity dynamic; PSD: power spectral density; SBI, simulation-based inference; MAF, masked autoregressive flow; NSF, neural spline flow; MCMC, Markov chain Monte Carlo.

*Corresponding author
Email addresses: `abolfazl.ziaee-mehr@univ-amu.fr` (Abolfazl Ziaeemehr), `viktor.jirsa@univ-amu.fr` (Viktor Jirsa [†]), `meysam.hashemi@univ-amu.fr` (Meysam Hashemi [†])
[†]These authors contributed equally.

1. Introduction

Understanding the complex dynamics of the brain and their neurobiological underpinnings, with the potential to advance precision medicine [1, 2, 3, 4], is a central goal in neuroscience. Modeling these dynamics provides crucial insights into causality and mechanisms underlying both normal brain function and various neurological disorders [5, 6, 7]. By integrating the average activity of large populations of neurons (e.g., neural mass models; [8, 9, 10, 11, 12, 13]) with information provided by structural imaging modalities (i.e., connectome; [14, 15, 16, 17]), whole-brain network modeling has proven to be a powerful tractable approach for simulating brain activities and emergent dynamics as recorded by functional imaging modalities (such as (s)EEG/MEG/fMRI; [18, 19, 20, 21, 22, 23].

The whole-brain models have been well-established in network neuroscience [24, 25] for understanding the brain structure and function [26, 27, 28, 29, 30, 31, 32] and investigating the mechanisms underlying brain dynamics at rest [33, 34, 35, 36], normal aging [37, 38], and also altered states such as anaesthesia and loss of consciousness [39, 40, 41, 42]. This class of computational models, also known as virtual brain models [43, 44, 18, 19, 45, 46], have shown remarkable capability in delineating the pathophysiological causes of a wide range of brain diseases, such as epilepsy [47, 48, 6], multiple sclerosis [49, 50], Alzheimer's disease [51, 52], Parkinson's disease [53, 54], neuropsychiatric disorders [55, 56], stroke [57, 58] and focal lesions [59]. In particular, they enable the personalized simulation of both normal and abnormal brain activities, along with their associated imaging recordings, thereby stratifying between healthy and diseased states [60, 61, 52] and potentially informing targeted interventions and treatment strategies [47, 62, 6, 45, 23]. Although there are only a few tools available for forward simulations at the whole-brain level, e.g., the brain network simulator The Virtual Brain (TVB; [44]), there is a lack of tools for addressing the inverse problem, i.e., finding the set of control (generative) parameters that best explains the observed data. This study aims to bridge this gap by addressing the inverse problem in large-scale brain networks, a crucial step toward making these models operable for clinical applications.

Accurately and reliably estimating the parameters of whole-brain models remains a formidable challenge, mainly due to the high-dimensionality and nonlinearity inherent in brain activity data, as well as the non-trivial effects of noise and network inputs. A large number of previous studies in whole-brain modeling have relied on optimization techniques to identify a single optimal value from an objective function, scoring the model's performance against observed data [34, 36, 63, 64]. This approach often involves minimizing metrics such as the Kolmogorov-Smirnov distance or maximizing the Pearson correlation between observed and generated data features such as functional connectivity (FC), functional connectivity dynamics (FCD), and/or power spectral density (PSD). Although fast, such a parametric approach results in only point estimates and fails to capture the relationship between parameters and their associated uncertainty. This limits the generalizability of findings, and hinders identifiability analysis, which explores the uniqueness of solutions. Furthermore, optimization algorithms can easily get stuck in local extrema, requiring multi-start strategies to address potential parameter degeneracies. These additional steps, while necessary, ultimately increase the computational cost. Critically, the estimation heavily depends on the form of the objective function defined for optimization [65, 66]. These limitations can be overcome by employing Bayesian inference, which naturally quantifies

the uncertainty in the estimation and statistical dependencies between parameters, leading to more robust and generalizable models. Bayesian inference is a principal method for updating prior beliefs with information provided by data through the likelihood function, resulting in a posterior probability distribution that encodes all the information necessary for inferences and predictions. This approach has proven essential for understanding the intricate relationships between brain structure and function [67, 37, 59], as well as for revealing the pathophysiological causes underlying brain disorders [68, 51, 49, 46, 69].

In this context, simulation-based inference (SBI; [70, 71, 68, 69]) has gained prominence as an efficient methodology for conducting Bayesian inference in complex models where traditional inference techniques become inapplicable. SBI leverages computational simulations to generate synthetic data and employs advanced probabilistic machine learning methods to infer the joint distribution over parameters that best explain the observed data, along with associated uncertainty. This approach is particularly well-suited for Bayesian inference on whole-brain models, which often exhibit complex dynamics that are difficult to retrieve from neuroimaging data with conventional estimation techniques. Crucially, SBI circumvents the need for explicit likelihood evaluation and the Markovian (sequential) property in sampling. Markov chain Monte Carlo (MCMC) is a standard non-parametric technique and asymptotically exact for sampling from a probability distribution [72]. However, for Bayesian inference on whole-brain models given high-dimensional data, the likelihood function becomes intractable, rendering MCMC sampling computationally prohibitive. SBI offers significant advantages, such as parallel simulation while leveraging amortized learning, making it effective for personalized inference from large datasets [69]. Amortization in artificial neural networks refers to the idea of reusing learned computations across multiple tasks or inputs [73]. Amortization in Bayesian inference refers to the process of training a shared inference network (e.g., a neural network) with with an intensive upfront computational cost, to perform fast inference across many different observations. Instead of re-running inference for each new observation, the trained model can rapidly return posterior estimates, significantly reducing computational cost at test time [68]. Following an initial computational cost during simulation and training to learn all posterior distributions, subsequent evaluation of new hypotheses can be conducted efficiently, without additional computational overhead for further simulations [68]. Importantly, SBI sidesteps the convergence issues caused by complex geometries that are often encountered when using gradient-based MCMC methods [74, 75, 76]. It also substantially outperforms approximate Bayesian computation (ABC) methods, which rely on a threshold to accept or reject samples [77, 78, 71]. Such a likelihood-free approach provides us with generic inference on complex systems as long as we can provide three modules:

(1) A prior distribution, describing the possible range of parameters from which random samples can be easily drawn, i.e., $\vec{\theta} \sim p(\vec{\theta})$.

(2) A simulator in computer code, that takes parameters as input and generates data as output, i.e., $\vec{x} \sim p(\vec{x} \mid \vec{\theta})$.

(3) A set of low-dimensional data features, which are informative of the parameters that we aim to infer.

These elements prepare us with a training data set $\{(\vec{\theta}_i, \vec{x}_i)\}_{i=1}^{N_{sim}}$ with a budget of N_{sim} simulations. Then, using a class of *deep neural density estimators* (such as masked autoregressive flows [79] or neural spline flows [80]), we can approximate

the posterior distribution of parameters given a set of observed data, i.e., $p(\vec{\theta} \mid \vec{x}_{obs})$. Therefore, a versatile toolkit should be flexible and integrative, adeptly incorporating these modules to enable efficient Bayesian inference over complex models.

To address the need for widely applicable, reliable, and efficient parameter estimation from different (source-localized) neuroimaging modalities, we introduce Virtual Brain Inference (`VBI`), a flexible and integrative toolkit for probabilistic inference at whole-brain level. This open-source toolkit offers fast simulation through just-in-time (JIT) compilation of various brain models in different programming languages (Python/C++) and devices (CPUs/GPUs). It supports space-efficient storage of simulated data (HDF5/NPZ/PT), provides a memory-efficient loader for batched data, and facilitates the extraction of low-dimensional data features (FC/FCD/PSD). Additionally, it enables the training of deep neural density estimators (MAFs/NSFs), making it a versatile tool for inference on neural sources corresponding to (s)EEG, MEG, fMRI recordings. `VBI` leverages high-performance computing, significantly enhancing computational efficiency through parallel processing of large-scale datasets, which would be impractical with current alternative methods. Although SBI has been used for low-dimensional parameter spaces [71, 49, 81], we demonstrate that it can scale to whole-brain models with high-dimensional unknown parameters, as long as informative data features are provided. `VBI` is now accessible on the cloud platform EBRAINS (https://ebrains.eu), enabling users to explore more realistic brain dynamics underlying brain (dys)functioning using Bayesian inference.

In the following sections, we will provide an overview of the theoretical foundations of whole-brain models and simulation-based inference, describe the architecture and features of the `VBI` toolkit, and demonstrate the validation through a series of case studies using in-silico data. We explore various whole-brain models corresponding to different types of brain recordings: a whole-brain network model of Wilson-Cowan [8], Jansen-Rit [82, 83], and Stuart-Landau [84] for simulating neural activity associated with EEG/MEG signals, the Epileptor [11] related to stereoelectro-EEG (sEEG) recordings, and Montbrió [12], and Wong-Wang [85, 86] mapped to fMRI BOLD signals. Although these models represent source signals and could be applied to other modalities (e.g., Stuart-Landau representing generic oscillatory dynamics), we focused on their capabilities to perform optimally in specific contexts. For instance, some are better suited for encephalographic signals (e.g., EEG/MEG) due to their ability to preserve spectral properties, while others have been used for fMRI data, emphasizing their ability to capture dynamic features such as bistability and time-varying functional connectivity.

2. Materials and methods

2.1. The virtual brain models

To build a virtual brain model (see Figure 1), the process begins with parcellating the brain into regions using anatomical data, typically derived from T1-MRI scans. Each region, represented as nodes in the network, is then equipped with a neural mass model to simulate the collective behavior of neurons within that area. These nodes are interconnected using a structural connectivity (SC) matrix, typically obtained from diffusion-weighted magnetic resonance imaging (DW-MRI). The connectome was built with TVB-specific reconstruction pipeline using generally available

neuroimaging software [16]. The entire network of interconnected nodes is then simulated using neuroinformatic tools, such as The Virtual Brain (`TVB`; [18]), generating neural activities at the source level. However, neural sources are not directly observable in real-world experiments, and a projection needs to be established to transform the simulated neural activity into empirically measurable quantities, such as (s)EEG, MEG, and fMRI. This approach offers insights into both normal brain function and neurological disorders [23]. In the following, we describe commonly used whole-brain network models at the source level, which can be mapped to different types of neuroimaging recordings. Note that each model represents one of many possible variants used in the literature, and the choice of model often depends on the specific research question, the spatial and temporal resolution of the available data, and the desired level of biological or mathematical detail.

2.1.1. Wilson-Cowan model

The Wilson-Cowan model [8] is a seminal neural mass model that describes the dynamics of connected excitatory and inhibitory neural populations, at cortical microcolumn level. It has been widely used to understand the collective behavior of neurons and simulate neural activities recorded by methods such as local field potentials (LFPs) and EEG. The model effectively captures phenomena such as oscillations, wave propagation, pattern formation in neural tissue, and responses to external stimuli, offering insights into various brain (dys)functions, particularly in Parkinson's disease [87, 88].

The Wilson-Cowan model describes the temporal evolution of the mean firing rates of excitatory (E) and inhibitory (I) populations using nonlinear differential equations. Each population's activity is governed by a balance of self-excitation, cross-inhibition, external inputs, and network interactions through long-range coupling. The nonlinearity arises from a sigmoidal transfer function $S_{i,e}(x)$, which maps the total synaptic input to the firing rate, capturing saturation effects and thresholds in neural response. In the whole-brain network extension, each neural population at node i receives input from other nodes via a weighted connectivity matrix, allowing the study of large-scale brain dynamics and spatial pattern formation [89, 90, 91].

$$
\begin{aligned}
\tau_e \frac{dE_k}{dt} &= -E_k + (k_e - r_e E_k) \cdot S_e \left(\alpha_e \left(c_{ee} E_k - c_{ei} I_k + P_k - \theta_e + g_e \sum_l \mathrm{SC}_{kl} E_l \right) \right) + \sigma \xi_k(t), \\
\tau_i \frac{dI_k}{dt} &= -I_k + (k_i - r_i I_k) \cdot S_i \left(\alpha_i \left(c_{ie} E_k - c_{ii} I_k + Q_k - \theta_i + g_i \sum_l \mathrm{SC}_{kl} I_l \right) \right) + \sigma \xi_k(t), \\
S_k(x) &= \begin{cases} c_k \left(\dfrac{1}{1 + e^{-a_k(x - b_k)}} - \dfrac{1}{1 + e^{a_k b_k}} \right), & \text{if shifted,} \\ \dfrac{c_k}{1 + e^{-a_k(x - b_k)}}, & \text{otherwise,} \quad k = i, e \end{cases}
\end{aligned}
\tag{1}
$$

which incorporates both local dynamics and global interactions, modulated by coupling strengths and synaptic weights. Here, SC_{kl} is an element of the (non)symmetric structural connectivity matrix and is nonzero if there is a connection between regions k and l. The nominal parameter values and the prior range for the target parameters are summarized in Table 1.

Parameter	Description	Value	Prior
c_{ee}	Excitatory to excitatory synaptic strength	16.0	
c_{ei}	Inhibitory to excitatory synaptic strength	12.0	
c_{ie}	Excitatory to inhibitory synaptic strength	15.0	
c_{ii}	Inhibitory to inhibitory synaptic strength	3.0	
τ_e	Time constant of excitatory population	8.0	
τ_i	Time constant of inhibitory population	8.0	
a_e	Sigmoid slope for excitatory population	1.3	
a_i	Sigmoid slope for inhibitory population	2.0	
b_e	Sigmoid threshold for excitatory population	4.0	
b_i	Sigmoid threshold for inhibitory population	3.7	
c_e	Maximum output of sigmoid for excitatory population	1.0	
c_i	Maximum output of sigmoid for inhibitory population	1.0	
θ_e	Firing threshold for excitatory population	0.0	
θ_i	Firing threshold for inhibitory population	0.0	
r_e	Refractoriness of excitatory population	1.0	
r_i	Refractoriness of inhibitory population	1.0	
k_e	Scaling constant for excitatory output	0.994	
k_i	Scaling constant for inhibitory output	0.999	
α_e	Gain of excitatory population	1.0	
α_i	Gain of inhibitory population	1.0	
P	External input to excitatory population	0.0	$\mathcal{U}(0,3)$
Q	External input to inhibitory population	0.0	$\mathcal{U}(0,3)$
g_e	Global coupling strength (excitatory)	0.0	$\mathcal{U}(0,1)$
g_i	Global coupling strength (inhibitory)	0.0	$\mathcal{U}(0,1)$
σ	Standard deviation of Gaussian noise	0.005	

Table 1. Parameter descriptions for capturing whole-brain dynamics using the Wilson-Cowan neural mass model.

2.1.2. Jansen-Rit model

The Jansen-Rit neural mass model [92] has been widely used to simulate physiological signals from various recording methods like intracranial LFPs, and scalp MEG/EEG recordings. For example, it has been shown to recreate responses similar to evoked-related potentials after a series of impulse stimulations [83, 93], generating high-alpha and low-beta oscillations (with added recurrent inhibitory connections and spike-rate modulation) [94], and also seizure patterns similar to those seen in temporal lobe epilepsy [95]. This biologically motivated model comprises three main populations of neurons: excitatory pyramidal neurons, inhibitory interneurons and excitatory interneurons. These populations interact with each other through synaptic connections, forming a feedback loop that produces oscillatory activity governed by a set of nonlinear ordinary differential equations [82, 83, 96]:

$$
\begin{aligned}
\dot{y}_{0i}(t) &= y_{3i}(t); \quad \dot{y}_{1i}(t) = y_{4i}(t); \quad \dot{y}_{2i}(t) = y_{5i}(t) \\
\dot{y}_{3i}(t) &= A\,a\,\mathrm{S}(y_{1i}(t) - y_{2i}(t)) - 2a\,y_{3i}(t) - a^2 y_{0i}(t) \\
\dot{y}_{4i}(t) &= A\,a\Big(P(t) + C_2\,\mathrm{S}(C_1 y_{0i}(t)) + G\,\mathrm{H}_i\Big) - 2a y_{4i}(t) - a^2 y_{1i}(t) \\
\dot{y}_{5i}(t) &= B\,b\Big(C_4\,\mathrm{S}(C_3 y_{0i}(t))\Big) - 2b\,y_{5i}(t) - b^2 y_{2i}(t) \\
\mathrm{S}(v) &= \frac{v_{max}}{1 + \exp(r(v_0 - v))} \\
\mathrm{H}_i &= \sum_{j=1}^{N} \mathrm{SC}_{ij}\,\mathrm{S}(y_{1j} - y_{2j})
\end{aligned}
\tag{2}
$$

Parameters	Description	Value	Prior
A	Excitatory PSPA	3.25 mV	
B	Inhibitory PSPA	22 mV	
$1/a$	Time constant of excitatory PSP	$a = 100\ s^{-1}$	
$1/b$	Time constant of inhibitory PSP	$b = 50\ s^{-1}$	
C_1, C_2	Average numbers of synapses between EP	$1C, 0.8C$	
C_3, C_4	Average numbers of synapses between IP	$0.25C$	
v_{max}	Maximum firing rate	5 Hz	
v_0	Potential at half of maximum firing rate	6 mV	
r	Slope of sigmoid function at v_0	$0.56\ mV^{-1}$	
C	Average numbers of synapses between neural populations	135	$\mathcal{U}(100, 500)$
G	Scaling the strength of network connections	1.5	$\mathcal{U}(0, 5)$

Table 2. Parameter descriptions for capturing whole-brain dynamics using Jansen-Rit neural mass model. EP: excitatory populations, IP: inhibitory populations, PSP: post synaptic potential, PSPA: post synaptic potential amplitude.

where y_{0i}, y_{1i}, and y_{2i} denote the average membrane potentials of pyramidal cells, excitatory interneurons, and inhibitory interneurons, respectively. Their corresponding time derivatives, $y_{3i}(t)$, $y_{4i}(t)$, and $y_{5i}(t)$, represent the rates of change of these membrane potentials. $P(t)$ also represents an external input current. The sigmoid function $S(x)$ maps the average membrane potential of neurons to their mean action potential firing rate. SC is a normalized structural connectivity matrix. The model's output at i_{th} region, corresponds to the membrane potential of pyramidal cells, and is given by $y_{1i} - y_{2i}$. The nominal parameter values and the prior range for the target parameters are summarized in Table 2.

2.1.3. Stuart-Landau oscillator

The Stuart-Landau oscillator [84] is a generic mathematical model used to describe oscillatory phenomena, particularly those near a Hopf bifurcation, which is often employed to study the nonlinear dynamics of neural activity [97, 98, 63, 49]. One approach uses this model to capture slow hemodynamic changes in BOLD signal [97], while others apply it to model fast neuronal dynamics, which can be linked directly to EEG/MEG data [98, 63, 49]. Note that this is a phenomenological framework, and both applications operate on completely different time scales.

In the network, each brain region, characterized by an autonomous Stuart-Landau oscillator, can exhibit either damped or limit-cycle oscillations depending on the bifurcation parameter a. If $a < 0$, the system shows damped oscillations, similar to a pendulum under friction. In this regime, the system, when subjected to perturbation, relaxes back to its stable fixed point through damped oscillations with an angular frequency ω_0. The rate of amplitude damping is determined by $|a|$. Conversely, if $a > 0$, the system supports limit cycle solutions, allowing for self-sustained oscillations even in the absence of external noise. At a critical value of $a = 0$, the system undergoes a Hopf bifurcation, i.e., small changes in parameters can lead to large variations in the system's behavior.

Using whole-brain network modeling of EEG/MEG data [49, 63], the oscillators are interconnected via white-matter pathways, with coupling strengths specified by subject-specific DTI fiber counts, i.e., elements of SC matrix. This adjacency matrix is then scaled by a global coupling parameter G. Note that coupling between regions accounts for finite conduction times, which are often estimated by dividing the Euclidean distances between nodes by an average conduction velocity $T_{jk} = d_{jk}/v$.

Parameters	Description	Value	Prior
a	Bifurcation parameter	-5	
ω_i	Natural angular frequency	$2\pi \cdot 40$ rad/s	
σ	Noise factor	10^{-4}	
v	Average conduction velocity	6.0 m/s	$\mathcal{U}(1, 30)$
G	Global coupling parameter	1200	$\mathcal{U}(0, 400)$

Table 3. Parameter descriptions for capturing whole-brain dynamics using Stuart-Landau oscillator.

Knowing the personalized time-delays [99, 100], we can use the distance as a proxy, assuming a constant propagation velocity. The distance itself can be defined as either the length of the tracts or the Euclidean distance. Taking this into account, the activity of each regions is given by a set of complex differential equations:

$$\dot{Z}_j = Z_j(a + i\omega_j - |Z_j|^2) + G\sum_{k=1}^{N} \mathrm{SC}_{jk}[Z_k(t - T_{jk}) - Z_j(t)] + \sigma(\xi_{j1}(t) + i\xi_{j2}(t)) \quad (3)$$

where Z is a complex variable, and $\mathrm{Re}[Z(t)]$ is the corresponding time series. In this particular realization, each region has a natural frequency of 40 Hz ($\omega_j = \omega_0 = 2\pi \cdot 40\ rad/s$), motivated by empirical studies demonstrating the emergence of gamma oscillations from the balance of excitation and inhibition, playing a role in local circuit computations [101].

In this study, for the sake of simplicity, a common cortico-cortical conduction velocity is estimated, i.e., the distance-dependent average velocities $v_{jk} = v$. We also consider $a = -5$, capturing the highly variable amplitude envelope of gamma oscillations as reported in experimental recordings [102, 63]. This choice also best reflects the slowest decay time constants of $\mathrm{GABA_B}$ inhibitory receptors–approximately 1 second [103]. A Gaussian noise (here denoted by ξ_i) with an intensity of $\sigma = 10^{-4}$ is added to each oscillator to mimic stochastic fluctuations. The nominal parameter values and the prior range for the target parameters are summarized in Table 3.

2.1.4. Epileptor model

In personalized whole-brain network modeling of epilepsy spread [47], the dynamics of each brain region are governed by the Epileptor model [11]. The Epileptor model provides a comprehensive description of epileptic seizures, encompassing the complete taxonomy of system bifurcations to simultaneously reproduce the dynamics of seizure onset, progression, and termination [104]. The full Epileptor model comprises five state variables that couple two oscillatory dynamical systems operating on three different time scales [11]. Then motivated by Synergetic theory [105, 106] and under time-scale separation [107], the fast variables rapidly collapse on the slow manifold [108], whose dynamics is governed by the slow variable. This adiabatic approximation yields the 2D reduction of whole-brain model of epilepsy spread, also known as the Virtual Epileptic Patient (VEP) as follows:

$$\begin{aligned} \dot{x}_i &= 1 - x_i^3 - 2x_i^2 - z_i + I_i \\ \dot{z}_i &= \frac{1}{\tau}(4(x_i - \eta_i) - z_i - G\sum_{j=1}^{N} \mathrm{SC}_{ij}(x_j - x_i)), \end{aligned} \quad (4)$$

where x_i and z_i indicate the fast and slow variables corresponding to i_{th} brain region, respectively, and the set of unknown η_i is the spatial map of epileptogenicity to be

estimated. SC is a normalized structure connectivity matrix. In real-world epilepsy applications [67, 68, 6], we compute the envelope function from sEEG data to perform inference. The nominal parameter values and the prior range for the target parameters are summarized in Table 4.

Parameter	Description	Value	Prior
I	Input electric current	3.1	
τ	System time constant	90 ms	
η_i	Spatial map of epileptogenicity	-3.65	$\mathcal{U}(-5,-1)$
G	Global scaling factor on network connections	1.0	$\mathcal{U}(0,2)$

Table 4. Parameter descriptions for capturing whole-brain dynamics using 2D Epileptor neural mass model.

2.1.5. Montbrió model

The exact macroscopic dynamics of a specific brain region (represented as a node in the network) can be analytically derived in thermodynamic limit of infinitely all-to-all coupled spiking neurons [12] or Θ neuron representation [109]. By assuming a Lorentzian distribution on excitabilities in large ensembles of quadratic integrate-and- fire neurons with synaptic weights J and a half-width Δ centered at η, the macroscopic dynamics has been derived in terms of the collective firing activity and mean membrane potential [12]. Then, by coupling the brain regions via an additive current (e.g., in the average membrane potential equations), the dynamics of the whole-brain network can be described as follow [59]:

$$\begin{aligned}\tau \dot{r_i}(t) &= 2r_i(t)v_i(t) + \frac{\Delta}{\pi\tau} \\ \tau \dot{v_i}(t) &= v_i^2(t) - (\pi\tau r_i(t))^2 + J\tau r_i(t) + \eta + \\ &\quad G\sum_{j=1}^{N} \mathrm{SC}_{ij} r_j(t) + I_{\mathrm{stim}}(t) + \xi(t),\end{aligned} \tag{5}$$

where v_i and r_i are the average membrane potential and firing rate, respectively, at the i_{th} brain region, and parameter G is the network scaling parameter that modulates the overall impact of brain connectivity on the state dynamics. The SC_{ij} denotes the connection weight between i_{th} and j_{th} regions, and the dynamical noise $\xi(t) \sim \mathcal{N}(0, \sigma^2)$ follows a Gaussian distribution with mean zero and variance σ^2.

The model parameters are tuned so that each decoupled node is in a bistable regime, exhibiting a down-state stable fixed point (low-firing rate) and an up-state stable focus (high-firing rate) in the phase-space [12, 81]. The bistability is a fundamental property of regional brain dynamics to ensure a switching behavior in the data (e.g., to generate FCD), that has been recognized as representative of realistic dynamics observed empirically [110, 111, 112, 59].

The solution of the coupled system yields a neuroelectric dataset that describes the evolution of the variables $(r_i(t), v_i(t))$ in each brain region i, providing measures of macroscopic activity. The surrogate BOLD activity for each region is then derived by filtering this activity through the Balloon-Windkessel model [113]. The input current I_{stim} represents the stimulation to selected brain regions, which increase the basin of attraction of the up-state in comparison to the down-state, while the fixed points move farther apart [110, 111, 112, 59].

The nominal parameter values and the prior range for the target parameters are summarized in Table 5.

Parameter	Description	Nominal value	Prior
τ	Characteristic time constant	1 ms	
J	Synaptic weight	14.5 ms^{-1}	
Δ	Spread of the heterogeneous noise distribution	0.7 ms^{-1}	
$I_{\text{stim}}(t)$	Input current representing stimulation	0.0	
σ	Gaussian noise variance	0.037	
η	Excitability	-4.6	$\mathcal{U}(-6, -3.5)$
G	Scaling the strength of network connections	0.56	$\mathcal{U}(0, 1)$

Table 5. Parameter descriptions for capturing whole-brain dynamics using Montbrió model.

2.1.6. Wong-Wang model

Another commonly used whole-brain model for simulation of neural activity is the so-called parameterized dynamics mean-field (pDMF) model [114, 36, 86]. At each region, it comprises a simplified system of two nonlinear coupled differential equations, motivated by the attractor network model, which integrates sensory information over time to make perceptual decisions, known as Wong-Wang model [85]. This biophysically realistic cortical network model of decision making then has been simplified further into a single-population model [86], which has been widely used to understand the mechanisms underpinning brain resting state dynamics [36, 115, 38]. The pDMF model has been also used to study whole-brain dynamics in various brain disorders, including Alzheimer's disease [116], schizophrenia [117], and stroke [118]. The pDMF model equations are given as:

$$\begin{aligned} \frac{dS_i(t)}{dt} &= -\frac{S_i}{\tau_s} + (1 - S_i)\gamma H(x_i) + \sigma\xi_i(t) \\ H(x_i) &= \frac{ax_i - b}{1 - \exp(-d(ax_i - b))} \\ x_i &= wJS_i + GJ\sum_{j=1}^{N} \text{SC}_{ij} S_j + I \end{aligned} \tag{6}$$

where $H(x_i)$ and S_i, and x_i denote the population firing rate, the average synaptic gating variable, and the total input current at the i_{th} brain region, respectively. $\xi_i(t)$ is uncorrelated standard Gaussian noise and the noise amplitude is controlled by σ. The nominal parameter values and the prior range for the target parameters are summarized in Table 6.

According to recent studies [36, 38], we can parameterize the set of w, I and σ as linear combinations of group-level T1w/T2w myelin maps [119] and the first principal gradient of functional connectivity:

$$\begin{aligned} w_i &= a_w \mathbf{Mye}_i + b_w \mathbf{Grad}_i + c_w \\ I_i &= a_I \mathbf{Mye}_i + b_I \mathbf{Grad}_i + c_I \\ \sigma_i &= a_\sigma \mathbf{Mye}_i + b_\sigma \mathbf{Grad}_i + c_\sigma \end{aligned} \tag{7}$$

where $\mathbf{Mye}_i$ and $\mathbf{Grad}_i$ are the average values of the T1w/T2w myelin map and the first FC principal gradient, respectively, within the i_{th} brain region. Therefore, the set of unknown parameters to estimate includes G and linear coefficients $(a_w, b_w, c_w, a_I, b_I, c_I, a_\sigma, b_\sigma, c_\sigma) \in \mathbb{R}^9$, hence 10 parameters in total.

Parameter	Description	Value	Prior
a	Max feeding rate of $H(x)$	270 n/C	
b	Half saturation of $H(x)$	108 Hz	
d	Control the steepness of curve of $H(x)$	0.154 s	
γ	Kinetic parameter	0.641/1000	
τ_s	Synaptic time constant	100 ms	
J	Synaptic coupling	0.2609 nA	
w	Local excitatory recurrence	0.6	$\mathcal{U}(0,1)$
I	Overall effective external input	0.3 nA	$\mathcal{U}(0,0.5)$
G	Scaling the strength of network connections	6.28	$\mathcal{U}(1,10)$
σ	Noise amplitude	0.005	$\mathcal{U}(0.0005,0.01)$

Table 6. Parameter descriptions for capturing whole-brain dynamics using the Wong-Wang model.

2.1.7. The Balloon-Windkessel model

The Balloon-Windkessel model is a biophysical framework that links neural activity to the BOLD signals detected in fMRI. This is not a neuronal model but rather a representation of neurovascular coupling, describing how neural activity influences hemodynamic responses. The model is characterized by two state variables: venous blood volume (v) and deoxyhemoglobin content (q). The system's input is blood flow (f_{in}), and the output is the BOLD signal (y):

$$\begin{aligned} y(t) &= \lambda(v,q,E_0) = V_0\big(k_1(1-q) + k_2(1-\frac{q}{v}) + k_3(1-v)\big) \\ k_1 &= 4.3\vartheta_0 E_0\,\mathrm{TE} \\ k_2 &= \varepsilon r_0 E_0\,\mathrm{TE} \\ k_3 &= 1-\varepsilon \end{aligned} \tag{8}$$

where V_0 represents the resting blood volume fraction, E_0 is the oxygen extraction fraction at rest, ϵ is the ratio of intra- to extravascular signals, r_0 is the slope of the relationship between the intravascular relaxation rate and oxygen saturation, ϑ_0 is the frequency offset at the surface of a fully deoxygenated vessel at 1.5 T, and TE is the echo time. The dynamics of venous blood volume v and deoxyhemoglobin content q are governed by the Balloon model's hemodynamic state equations:

$$\begin{aligned} \tau_0\frac{dv}{dt} &= f(t) - v(t)^{1/\alpha} \\ \tau_0\frac{dq}{dt} &= f(t)\frac{1-(1-E_0)^{1/f}}{E_0} - v(t)^{1/\alpha}q(t) \end{aligned} \tag{9}$$

where τ_0 is the transit time of blood flow, α reflects the resistance of the venous vessel (stiffness), and $f(t)$ denotes blood inflow at time t, given by

$$\frac{df}{dt} = s,$$

where s is an exponentially decaying vasodilatory signal defined by

$$\frac{ds}{dt} = \epsilon x - \frac{s}{\tau_s} - \frac{(f-1)}{\tau_f} \tag{10}$$

where, ϵ represents the efficacy of neuronal activity $x(t)$ (i.e., integrated synaptic activity) in generating a signal increase, τ_s is the time constant for signal decay, and τ_f is the time constant for autoregulatory blood flow feedback [113]. For parameter values, see Table 7, taken from [113, 120, 121]. The resulting time series is downsampled to match the TR value in seconds.

Parameter	Description	Value
τ_s	Rate constant of vasodilatory signal decay in seconds	1.5
τ_f	Time of flow-dependent elimination in seconds	4.5
α	Grubb's vessel stiffness exponent	0.2
τ_0	Hemodynamic transit time in seconds	1.0
ϵ	Efficacy of synaptic activity to induce signal	0.1
r_0	Slope of intravascular relaxation rate in Hertz	25.0
ϑ_0	Frequency offset at outer surface of magnetized vessels	40.3
ε	Ratio of intra- and extravascular BOLD signal at rest	1.43
V_0	Resting blood volume fraction	0.02
E_0	Resting oxygen extraction fraction	0.8
TE	Echo time, 1.5T scanner	0.04

Table 7. Parameter descriptions for the Balloon-Windkessel model to map neural activity to the BOLD signals detected in fMRI.

2.2. Simulation-based inference

In the Bayesian framework [122], parameter estimation involves quantifying and propagating uncertainty through probability distributions placed on the parameters (prior information before seeing data), which are updated with information provided by the data (likelihood function). The formidable challenge to conducting efficient Bayesian inference is evaluating the likelihood function $p(x \mid \theta)$. This typically involves intractable integrating over all possible trajectories in the latent space: $p(x \mid \theta) = \int p(x, z \mid \theta) dz$, where $p(x, z \mid \theta)$ is the joint probability density of the data x and latent variables z, given parameters θ. For whole-brain network models with high-dimensional and nonlinear latent spaces, the computational cost can be prohibitive, making likelihood-based inference with MCMC sampling challenging to converge [76, 67, 123].

Simulation-based inference (SBI; [70, 71, 68]), or likelihood-free inference [124, 125, 126], addresses issues with explicit likelihood evaluation in complex systems, where the it often becomes intractable. The task of density estimation, one of the most fundamental problems in statistics, is to infer an underlying probability distribution based on a set of independently and identically distributed data points drawn from that distribution. Traditional density estimators, such as histograms and kernel density estimators, typically perform well only in low-dimensional settings. Recently, neural network-based approaches have been proposed for conditional density estimation, showing promising results in Bayesian inference problems involving high-dimensional data [127, 79, 125, 71, 128, 129, 68].

Given a prior distribution $p(\vec{\theta})$ placed on the parameters $\vec{\theta}$, N random simulations are generated (with samples from prior), resulting in pairs $\{(\vec{\theta}_i, \vec{x}_i)\}_{i=1}^{N_{sim}}$, where $\vec{\theta}_i \sim p(\vec{\theta})$ and $\vec{x}_i \sim p(\vec{x} \mid \vec{\theta})$ is the simulated data given $\vec{\theta}_i$. By training a deep neural density estimator F (such as normalizing flows [79, 80]), we can approximate the posterior $p(\vec{\theta} \mid \vec{x})$ with $q_{F(\phi,\vec{x})}$ by minimizing the loss function

$$\mathcal{L}(\phi) = -\sum_{i=1}^{N_{sim}} \log q_{F(\phi,x_i)}(\theta_i) \tag{11}$$

over network parameters ϕ. Once the parameters of the neural network F are optimized, for observed data $\vec{x}_{obs}$ we can readily estimate the target posterior $q_{F(\phi,\vec{x_{obs}})}(\vec{\theta}) \simeq p(\vec{\theta} \mid \vec{x}_{obs})$. This allows for rapid approximation and sampling from the posterior dis-

tribution for any new observed data through a forward pass in the trained network [68, 69].

This approach uses a class of generative machine learning models known as normalizing flows ([130, 131]) to transform a simple based distribution into any complex target through a sequence of invertible mappings. Here, generative modeling is an unsupervised machine learning method for modeling a probability distribution given samples drawn from that distribution. The state-of-the-art normalizing flows, such as masked autoregressive flows (MAF; [79]) and neural spline flows (NSF; [80]), enable fast and exact density estimation and sampling from high-dimensional distributions. These models learn mappings between input data and probability densities, capturing complex dependencies and multi-modal distributions [131, 132].

In our work, we integrate the implementation of these models from the open-source `SBI` tool, leveraging both MAF and NSF architectures. The MAF model comprises 5 flow transforms, each with two blocks and 50 hidden units, tanh nonlinearity and batch normalization after each layer. The NSF model consists of 5 flow transforms, two residual blocks of 50 hidden units each, ReLU nonlinearity, and 10 spline bins. We apply these generative models to virtual brain simulations conducted with random parameters, to approximate the full posterior distribution of parameters from low-dimensional data features. Note that we employ a single round of SBI to benefit from amortization strategy rather than using a sequential approach that is designed to achieve a better fit but only for a specific dataset [68, 69].

2.3. Sensitivity analysis

Sensitivity analysis is a crucial step for identifying which model parameters influence the model's behavior in response to changes in input [66, 68]. A local sensitivity can be quantified by computing the curvature of objective function through the Hessian matrix [133, 66]. Using SBI, after estimating the posterior for a specific observation, we can perform sensitivity analysis by computing the eigenvalues and corresponding eigenvectors of the following matrix [134, 135]:

$$M = \mathbb{E}_{p(\theta|x_{obs})}[\nabla_\theta p(\theta|x_{obs})^T \nabla_\theta p(\theta|x_{obs})], \tag{12}$$

which then does an eigendecomposition $M = Q\Lambda Q^{-1}$. A large eigenvalue in the so-called active subspaces [136] indicates that the gradient of the posterior is large in the corresponding direction, suggesting that the system output is sensitive to changes along that eigenvector.

2.4. Evaluation of posterior fit

To assess the reliability of Bayesian inference using synthetic data, we evaluate the posterior z-scores (denoted by z) against the posterior shrinkage (denoted by s), as defined by [75]:

$$z = |\frac{\bar{\theta} - \theta^*}{\sigma_{post}}|, \tag{13}$$

$$s = 1 - \frac{\sigma_{post}^2}{\sigma_{prior}^2}, \tag{14}$$

where $\bar{\theta}$ and θ^* are the posterior mean and the true values, respectively, σ_{prior} is the standard deviation of the prior, and σ_{post} is the standard deviation of the posterior.

The z-score quantifies how far the posterior mean of a parameter lies from a reference value (e.g., the true value), scaled by the posterior standard deviation. The shrinkage quantifies how much the posterior distribution has contracted relative to the initial prior distribution after learning from data. A small z-score indicates that the posterior estimate is close to the true value, reflecting accurate inference. A large shrinkage value suggests that the posterior is sharply concentrated, indicating that the parameter is well-identified. According to these definitions, an ideal Bayesian inference is characterized by z-scores close to zero and posterior shrinkage values close to one, reflecting both accuracy and reliability in the inferred parameters.

2.5. Flexible simulator and model building

A key feature of the `VBI` pipeline is its modularity and flexibility in integrating various simulators (see Figure 2). The *Simulation* module of the `VBI` pipeline is designed to be easily interchangeable, allowing researchers to replace it with other simulators, such as `TVB` [18], `Neurolib` [137], `Brian` [138], `Brainpy` [139]. This adaptability supports a wide range of simulation needs and computational environments, making the `VBI` a versatile tool for inference in system neuroscience. In particular, *Simulation* module offers a comprehensive implementation of commonly used whole-brain models. This is a customized version of implementation from open-source `TVB` simulator. While `VBI` does not encompass all the features of the original `TVB`, it is mainly designed to leverage the computational power of GPU devices and significantly reduce RAM requirements (see Figure S1**A**). This optimization ensures that high-performance clusters can be fully utilized, enabling parallel and scalable simulations, as often is required to perform scalable SBI.

2.6. Comprehensive feature extraction

`VBI` offers a comprehensive toolbox for feature extraction across various datasets. The *Features* module includes but not limited to: (1) *Statistical features*, including mean (average of elements), variance (spread of the elements around mean), kurtosis (tailedness of the distribution of elements), and skewness (the asymmetry of the distribution of elements), that can be applied to any matrix. (2) *Spectral features*, such as low-dimensional summary statistics of power spectrum density (PSD). (3) *Temporal features*, including zero crossings, area under the curve, average power, and envelope. (4) *Connectivity features*, including functional connectivity (FC), which represents the statistical dependencies or correlations between activity patterns of different brain regions, and functional connectivity dynamics (FCD), which captures the temporal variations and transitions in these connectivity patterns over time. These calculations are performed for the whole-brain and/or subnetwork (e.g., limbic system, resting state networks). However, since these matrices are still high-dimensional, we use standard dimensional reduction techniques, such as principal component analysis (PCA) on FC/FCD matrices, to extract their associated low-dimensional summary statistics. (5) *Information theory features*, such as mutual information and transfer entropy. Following [69], we use the term *spatio-temporal data features* to refer to both *statistical features* and *temporal features* derived from time series. In contrast, we refer to the *connectivity features* extracted from FC/FCD matrices as *functional data features*.

The *Features* module uses parallel multiprocessing to speed up feature calculation. Additionally, it provides flexibility for users to add their own custom feature

calculations with minimal effort and expertise, or to adjust the parameters of existing features based on the type of input time series. The feature extraction module is designed to be interchangeable with existing feature extraction libraries such as `tsfel` [140], `pyspi` [141], `hctsa`, [142], and `scikit-learn` [143]. Note that some lightweight libraries such as `catch22` [144] are directly accessible from the `VBI` feature extraction module.

2.7. VBI workflow

Figure 1 illustrates an overview of our approach in `VBI`, which combines virtual brain models and simulation-based inference to make probabilistic predictions on brain dynamics from (source-localized) neuroimaging recordings. The inputs to the pipeline include the structural imaging data (for building the connectome), functional imaging data such as (s)EEG/MEG, and fMRI as the target for fitting, and prior information as a plausible range over control parameters for generating random simulations. The main computational costs involve model simulations and data feature extraction. The output of the pipeline is the joint posterior distribution of control parameters (such as excitability, synaptic weights, or effective external input) that best explains the observed data. Since the approach is amortized (i.e., it learns across all combinations in the parameter space), it can be readily applied to any new data from a specific subject.

In the first step, non-invasive brain imaging data, such as T1-weighted MRI and Diffusion-weighted MRI (DW-MRI), are collected for a specific subject (Figure 1**A**). T1-weighted MRI images are processed to obtain brain parcellation, while DW-MRI images are used for tractography. Using the estimated fiber tracts and the defined brain regions from the parcellation, the connectome (i.e., the complete set of links between brain regions) is constructed by counting the fibers connecting all regions. The SC matrix, with entries representing the connection strength between brain regions, forms the structural component of the virtual brain which constrains the generation of brain dynamics and functional data at arbitrary brain locations (e.g., cortical and subcortical structures).

Subsequently, each brain network node is equipped with a computational model of average neuronal activity, known as neural mass models (see Figure 1**B** and subsection 2.1). They can be represented in the generic form of a dynamical model as $\dot{\vec{x}} = f(\vec{x}, \vec{\theta}, I_{input})$, with the system variables $\vec{x}$ (such as membrane potential and firing rate), the control parameters $\vec{\theta}$ (such as excitability), and the input current I_{input} (such as stimulation). This integration of mathematical mean-field modeling (neural mass models) with anatomical information (connectome) allows us to efficiently analyze functional neuroimaging modalities at the whole-brain level.

To quantify the posterior distribution of control parameters given a set of observations, $p(\vec{\theta} \mid \vec{x})$, we first need to define a plausible range for the control parameters based on background knowledge $p(\vec{\theta})$, i.e., a simple base distribution known as a prior. We draw random samples from the prior and provide them as input to the `VBI` simulator (implemented by *Simulation* module) to generate simulated time series associated with neuroimaging recordings, as shown in Figure 1**C**. Subsequently, we extract low-dimensional data features (implemented by *Features* module), as shown in Figure 1**D** for FC/FCD/PSD, to prepare the training dataset $\{(\vec{\theta}_i, \vec{x}_i)\}_{i=1}^{N_{sim}}$, with a budget of N_{sim} simulations. Then, we use a class of deep neural density estimators, such as MAF or NSF models, as schematically shown in Figure 1**E**, to learn all the

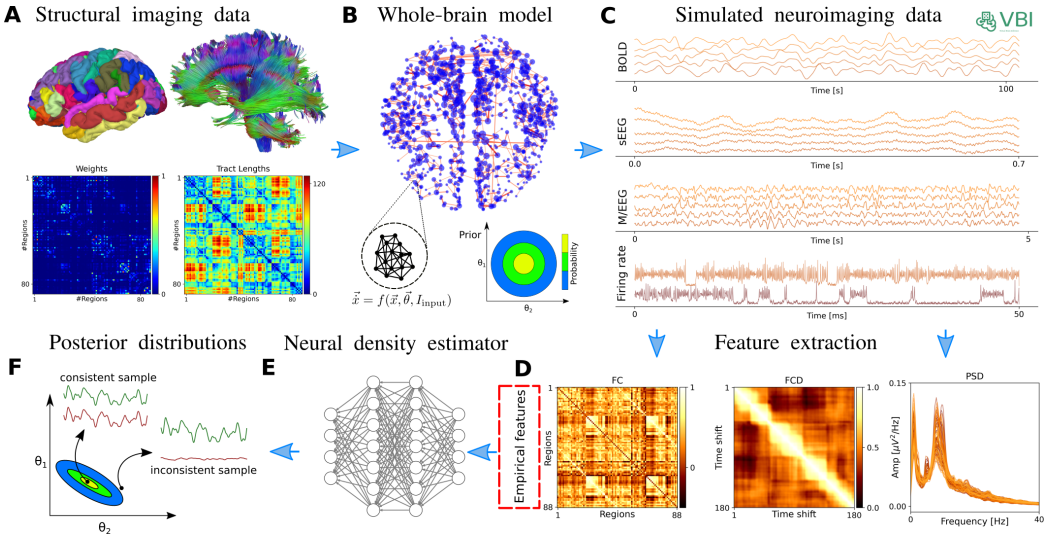

Fig. 1. The workflow of Virtual Brain Inference (`VBI`). This probabilistic approach is designed to estimate the posterior distribution of control parameters in virtual brain models from whole-brain recordings. (**A**) The process begins with constructing a personalized connectome using diffusion tensor imaging and a brain parcellation atlas. (**B**) The personalized virtual brain model is then assembled. Neural mass models describing the averaged activity of neural populations, in the generic form of $\dot{\vec{x}} = f(\vec{x}, \vec{\theta}, I_{input})$, are placed to each brain region, and interconnected via the structural connectivity matrix. Initially, the control parameters are randomly drawn from a simple prior distribution. (**C**) Next, the `VBI` operates as a simulator that uses these samples to generate time series data associated with neuroimaging recordings. (**D**) We extract a set summary statistics from the low-dimensional features of the simulations (FC, FCD, PSD) for training. (**E**) Subsequently, a class of deep neural density estimators is trained on pairs of random parameters and their corresponding data features to learn the joint posterior distribution of the model parameters. (**F**) Finally, the amortized network allows us to quickly approximate the posterior distribution for new (empirical) data features, enabling us to make probabilistic predictions that are consistent with the observed data.

posterior $p(\vec{\theta} \mid \vec{x})$. Finally, we can readily sample from $p(\vec{\theta} \mid \vec{x_{obs}})$, which determines the probability distribution in parameter space that best explains the observed data.

Figure 2 depicts the structure of the `VBI` toolkit, which consists of three main modules. The first module, referred to as *Simulation* module, is designed for fast simulation of whole-brain models, such as Wilson-Cowan (subsubsection 2.1.1), Jansen-Rit (subsubsection 2.1.2), Stuart-Landau (subsubsection 2.1.3), Epileptor (subsubsection 2.1.4), Montbrió (subsubsection 2.1.5), and Wong-Wang (subsubsection 2.1.6). These whole-brain models are implemented across various numerical computing libraries such as Cupy (GPU-accelerated computing with Python), C++ (a high-performance systems programming language), Numba (a just-in-time compiler for accelerating Python code), and PyTorch (an open-source machine learning library for creating deep neural network).

The second module, *Features*, provides a versatile tool for extracting low-dimensional features from simulated time series (see subsection 2.6). The features include, but are not limited to, *spectral*, *temporal*, *connectivity*, *statistical*, and *information theory* related features, and the associated summary statistics. The third module focuses on *Inference*, i.e., training the deep neural density estimators, such as MAF and NSF (see subsection 2.2), to learn the joint posterior distribution of control parameters.

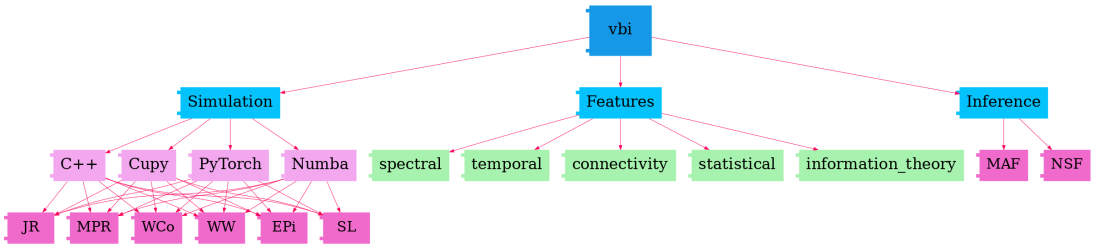

Figure 2. Flowchart of the `VBI` Structure. This toolkit consists of three main modules: (1) The *Simulation* module, implementing various whole-brain models, such as Wilson-Cowan (WCo), Jansen-Rit (JR), Stuart-Landau (SL), Epileptor (EPi), Montbrió (MPR), and Wong-Wang (WW), across different numerical computing libraries (C++, Cupy, PyTorch, Numba). (2) The *Features* module, offering an extensive toolbox for extracting low-dimensional data features, such as spectral, temporal, connectivity, statistical, and information theory features. (3) The *Inference* module, providing neural density estimators (such as MAF and NSF) to approximate the posterior of parameters.

3. Results

In the following, we demonstrate the capability of VBI for inference on the state-of-the-art whole-brain network models using in-silico testing, where the ground truth is known. We apply this approach to simulate neural activity and associated measurements, including (s)EEG/MEG, and fMRI, while also providing diagnostics for the accuracy and reliability of the estimation. Note that for (s)EEG/MEG neuroimaging, we perform inference at the regional level rather than at the sensor level, whereas for fMRI, it is mapped using the Balloon-Windkessel model (see subsubsection 2.1.7). The results presented are based on synthetic data generated using a set of predefined parameters, referred to as the ground truth, randomly selected within biologically plausible ranges and incorporating a certain level of heterogeneity.

We first demonstrate inference on the whole-brain network model of the Wilson-Cowan (Eq. 1), which is capable of generating a wide range of oscillatory dynamics depending on the control parameters. Specifically, we estimate the bifurcation parameters $P_i \in \mathbb{R}^{88}$, representing the external input to each excitatory population, and the global excitatory coupling parameter g_e. Figure 3**A** and **B** present the observed and predicted EEG-like signals, represented by the activity of excitatory populations E across regions. Panels **C** and **D** show the corresponding power spectral density (PSD), as data features. Panels **E** and **F** illustrate the inferred posterior distributions for parameters P_i and g_e, respectively, given $\theta = \{g_e, P_i\} \in \mathbb{R}^{89}$. For training, we conducted 250k random simulations from uniform priors $g_e \sim \mathcal{U}(0, 3)$ and $P_i \sim \mathcal{U}(0, 3)$ (shown in blue; see Table 1). After approximately 2 hours of training using MAF density estimators, posterior sampling was completed within a few seconds. Due to the large number of simulations and informativeness of data features, we achieved accurate estimations of the high-dimensional and heterogeneous control parameters. Ground-truth values (shown in green) are well recovered, leading to close agreement between observed and predicted PSDs of the signals. Finally, Figure 3**G** reports the posterior shrinkage and z-score metrics used to evaluate the quality of the parameter estimation. The results indicate that the inferred posteriors are both precise and well-centered around the ground-truth values, as reflected by high shrinkage and low z-scores. See Figure S2 for estimation over other configurations. Moreover, Figure S3, and Figure S4, indicate the higher accuracy of NSF, though with substantially more

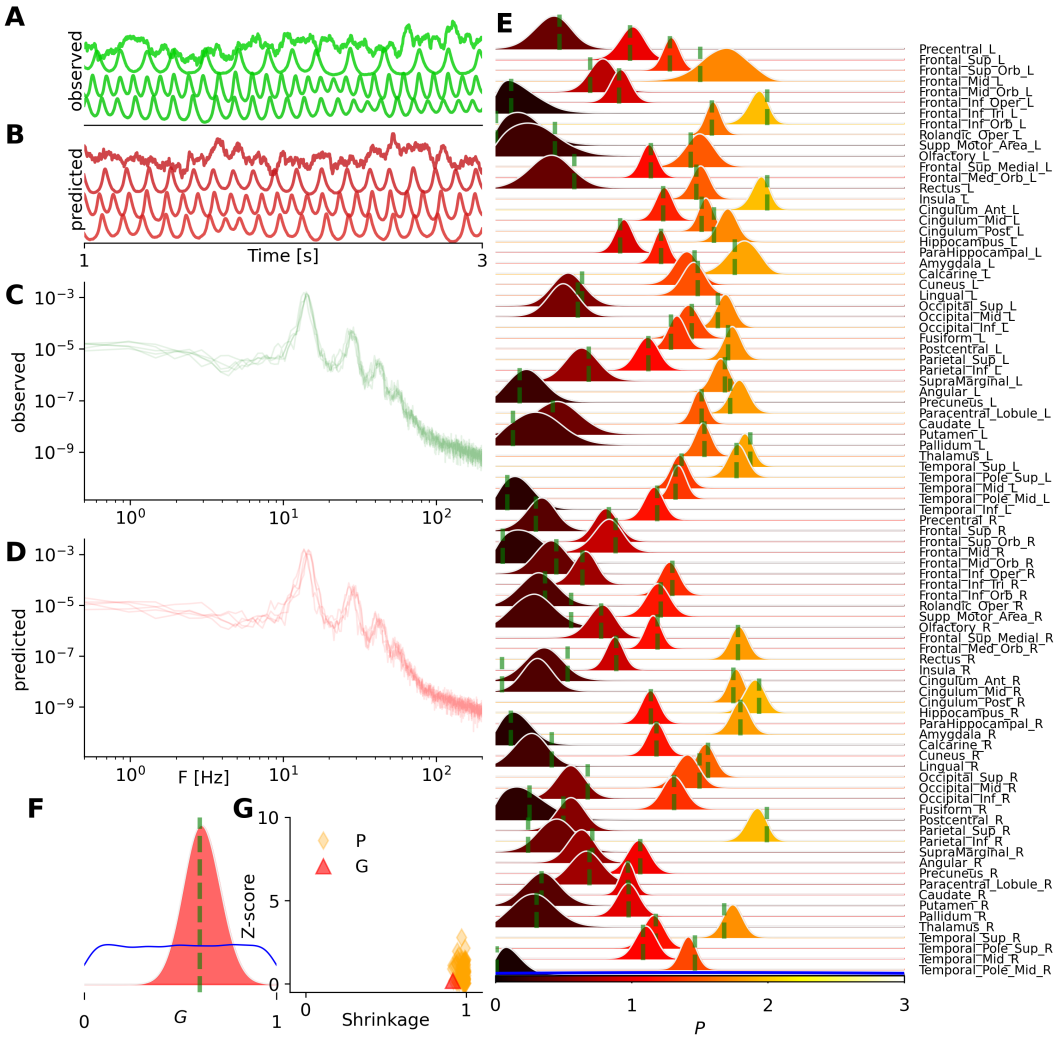

Figure 3. Bayesian inference on heterogeneous control parameters in the whole-brain network of Wilson-Cowan model. The set of inferred parameters is $\vec{\theta} = \{G, P_i\} \in \mathbb{R}^{89}$, with the global scaling parameter G and average external input current to excitatory populations per P_i region, given $i \in \{1, 2, ..., N_n = 88\}$ parcelled regions. Summary statistics of power spectrum density (PSD) were used for training of MAF density estimator, with a budget of 260k simulations. (**A**) and (**B**) illustrate the observed and predicted neural activities, respectively. (**C**) and (**D**) show the observed and predicted PSDs as the data. (**E**) and (**F**) display the posterior distribution of P_i per region, and global coupling G, respectively. The ground truth and prior are represented by a vertical green line and a blue distribution, respectively. (**G**) shows the inference evaluation using posterior shrinkage and z-score.

computational cost for training compared to MAF.

Then, we demonstrate the inference on heterogeneous control parameters in the whole-brain network of Jansen-Rit (see subsubsection 2.1.2), commonly used for modeling EEG/MEG data, e.g., in dementia and Alzheimer's disease [145, 146]. Figure 4**A** and **B** show the observed and predicted EEG signals, given by $(y_{1i} - y_{2i})$ at each region, while Figure 4**C** and **D** illustrate the observed and predicted features such as PSD, respectively. Figure 4**E** and **F** show the estimated posterior distributions of synaptic connections C_i, and the global coupling parameter G, respectively, given the set of unknown parameters $\vec{\theta} = \{G, C_i\} \in \mathbb{R}^{89}$. Here we conducted

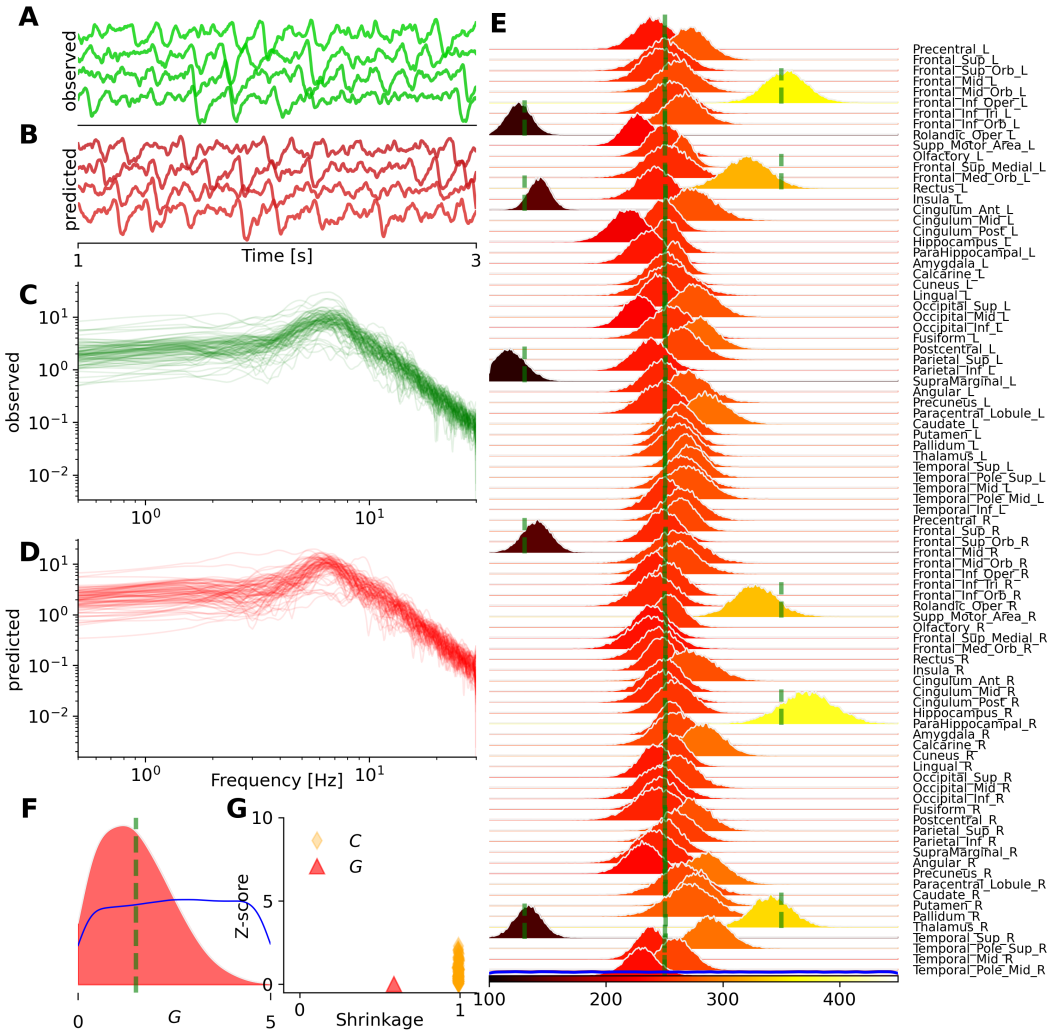

Figure 4. Bayesian inference on heterogeneous control parameters in the whole-brain network of Jansen-Rit model. The set of inferred parameters is $\vec{\theta} = \{G, C_i\} \in \mathbb{R}^{89}$, with the global scaling parameter G and average numbers of synapse between neural populations per C_i region, given $i \in \{1, 2, ..., N_n = 88\}$ parcelled regions. Summary statistics of power spectrum density (PSD) were used for training, with a budget of 50k simulations. (**A**) and (**B**) illustrate the observed and predicted neural activities, respectively. (**C**) and (**D**) show the observed and predicted data features, such as PSD. (**E**) and (**F**) display the posterior distribution of C_i per region, and global coupling G, respectively. The ground truth and prior are represented by vertical green lines and a blue distribution, respectively. (**G**) shows the inference evaluation using the shrinkage and z-score of the estimated posterior distributions.

50k random simulations with samples drawn from uniform priors $G \in \mathcal{U}(0, 5)$ and $C_i \in \mathcal{U}(100, 650)$ (shown in blue, see Table 2). After approximately 45 min of training (MAF density estimator), the posterior sampling took only a few seconds. With such a sufficient number of simulations and informative data features, `VBI` shows accurate estimation of high-dimensional heterogeneous parameters (given the ground truth, shown in green), leading to a strong correspondence between the observed and predicted PSD of EEG/MEG data. Figure 4**G** displays the shrinkage and z-score as the evaluation metrics, indicating an ideal Bayesian estimation for C_i parameters, but not for the coupling parameter G. This occurred because the network input did not

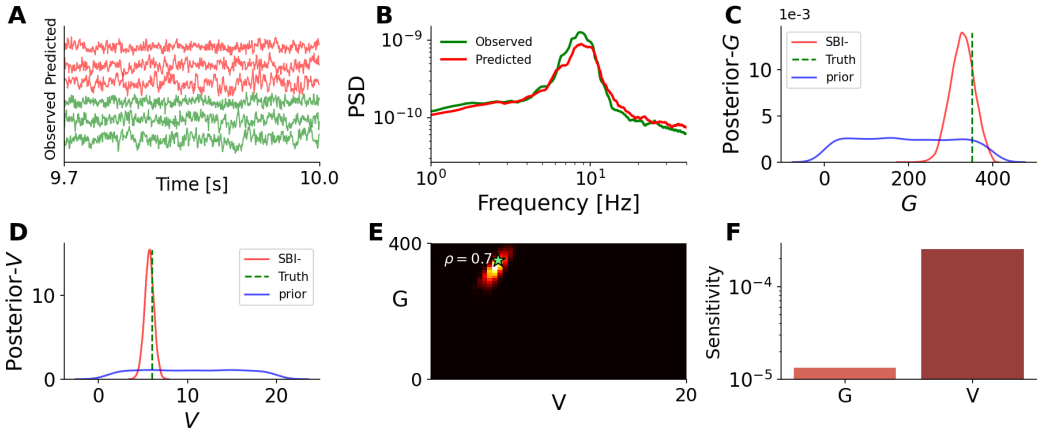

Figure 5. Bayesian inference on global scaling parameter G and the averaged velocity V of signal transmission using the whole-brain network model of Stuart-Landau oscillators. The set of estimated parameters is $\vec{\theta} = \{G, V\} \in \mathbb{R}^2$, and the summary statistics of PSD signals with a budget of 2k simulations were used for training. (**A**) illustrates exemplary observed and predicted neural activities (in green and red, respectively). (**B**) shows the observed and predicted PSD signals (in green and red, respectively). (**C**) and (**D**) display the posterior distribution of averaged velocity V and global coupling G, respectively. The true values and prior are shown as vertical green lines and a blue distribution, respectively. (**E**) shows the joint posterior distribution indicating a high correlation between posterior samples. (**F**) illustrates the sensitivity analysis based on the eigenvalues of the posterior distribution.

induce a significant change in the intrinsic frequency of activities at regional level, resulting in diffuse uncertainty in its estimation for this model.

Note that relying on only alpha-peak while excluding other summary statistics, such as total power (i.e., area under the curve), leads to poor estimation of synaptic connections across brain regions (see Figure S5). This results in less accurate predictions of the PSD, with more dispersion in their amplitudes. This example demonstrates that `VBI` provides a valuable tool for hypothesis evaluation and improved insight into data features by uncertainty quantification, and their impact on predictions.

To demonstrate efficient inference on the whole-brain time delay from EEG/MEG data, we used a whole-brain network model of coupled generic oscillators (Stuart-Landau model (see subsubsection 2.1.3). This model could establish a causal link between empirical spectral changes and the slower conduction velocities observed in multiple sclerosis patients, resulting from immune system attacks on the myelin sheath [49, 50]. The parameter set to estimate is $\vec{\theta} = \{G, V\} \in \mathbb{R}^2$, consisting of the global scaling parameter G and the averaged velocity of signal transmission V. The training was performed using a budget of only 2k simulations, which was sufficient due to the low dimensionality of the parameter space. Figure 5**A** illustrates the comparison between observed (in green) and predicted neural activities (in red). Figure 5**B** shows a close agreement between observed and predicted PSD signals, as the data feature used for training. Figure 5**C** and **D** provide visualizations of the posterior distributions for the averaged velocity V and the global coupling G. In these panels, we can see a large shrinkage in the posterior (in red) the uniform prior (in blue) centered around the true values (vertical green lines). Importantly, Figure 5**E** presenting the joint posterior distribution, indicates a high correlation of $\rho = 0.7$ between parameters G and V. This illustrates the advantage of Bayesian

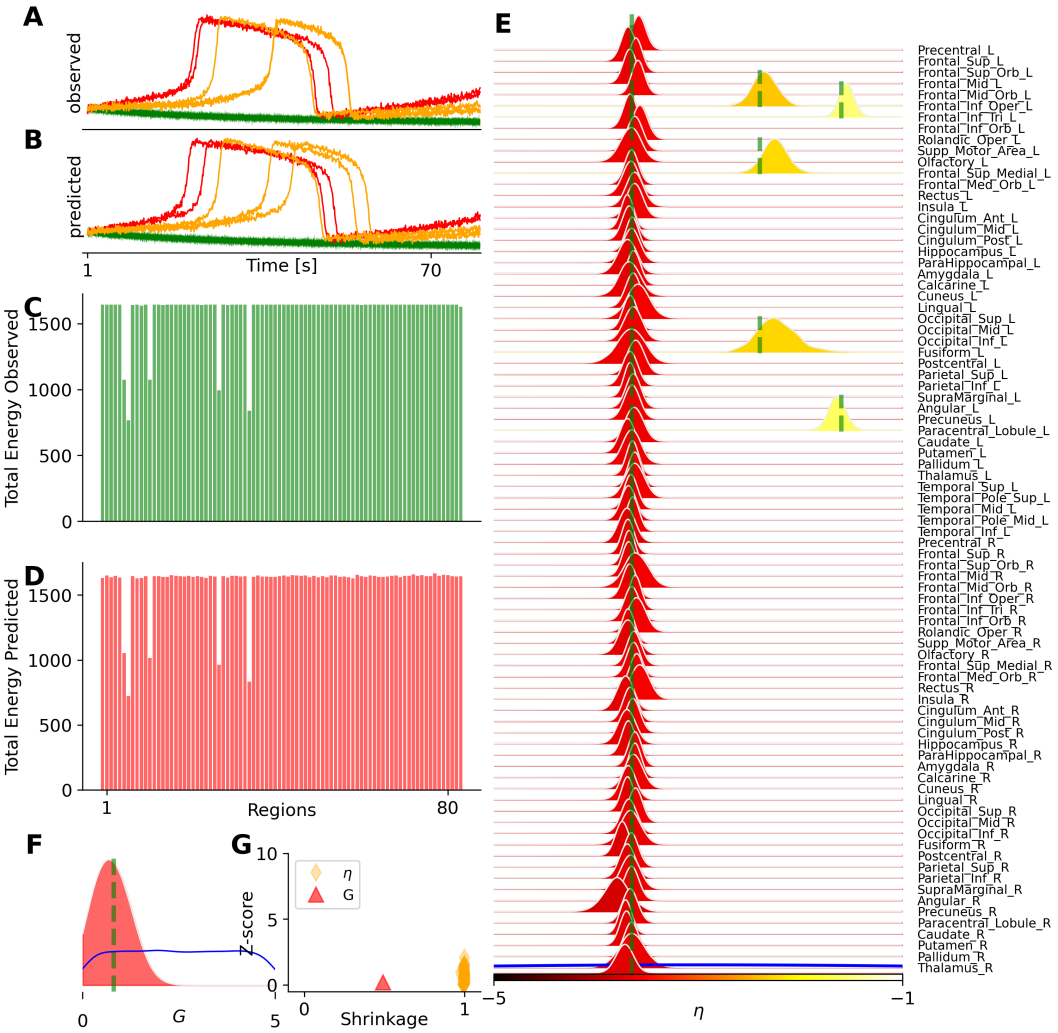

Figure 6. Bayesian inference on the spatial map of epileptogenicity across brain regions in the VEP model. The set of inferred parameters is $\vec{\theta} = \{G, \eta_i\} \in \mathbb{R}^{89}$, as the global scaling parameter and spatial map of epileptogenicity with $i \in \{1, 2, ..., N_n = 88\}$ parcelled regions. (**A**) The observed seizure envelope generated by the Epileptor model, given two regions as epileptogenic zones (in red) and three regions as propagation zones (in yellow), while the rest are healthy (in green). (**B**) The predicted seizure envelope, by training MAF model on a dataset containing 10k simulations, using only the total power and seizure onset per region as the data features. (**C**) and (**D**) show the observed and predicted data features, respectively. (**E**) and (**F**) show the posterior distributions of heterogeneous control parameters η_i, and global coupling parameter G, respectively. (**G**) The posterior z-scores versus posterior shrinkages for estimated parameters.

estimation in identifying statistical relationships between parameters, which helps to detect degeneracy among them. This is crucial for causal hypothesis evaluation and guiding conclusions in clinical settings. Finally, Figure 5**F** illustrates the sensitivity analysis (based on the eigenvalues of the posterior distribution), revealing that the posterior is more sensitive to changes in V compared to G. This highlights the relative impact of these parameters on the model's posterior estimates.

Next, we demonstrate the inference on a whole-brain model of epilepsy spread, known as the Virtual Epileptic Patient (VEP; [47, 76]), used to delineate the epileptogenic and propagation zone networks from the invasive sEEG recordings (subsub-

section 2.1.4). Here, we used a large value for system time constant $\tau = 90\ ms$ (see Table 4) to generate slow-fast dynamics in pathological areas, corresponding to seizure envelope at each brain region. Figure 6 demonstrates the inference the set of inferred parameters $\vec{\theta} = \{G, C_i\} \in \mathbb{R}^{89}$, with the global scaling parameter G and spatial map of epileptogenicity η_i, given $i \in \{1, 2, ..., N_n = 88\}$ parcelled regions. Figure 6**A** and **B** show the observed and predicted envelope, respectively, at each brain region. Here, the whole brain regions are classified into two epileptogenic zones (in red, corresponding to high excitability), three propagation zones (in yellow, corresponding to excitability close to bifurcation), and the rest as healthy regions (in green, corresponding to low excitability). Figure 6**C** and **D** illustrate the observed and predicted data features as the total power energy per region, calculated as the area under the curve. Additionally the seizure onset at each region was used as a data feature for training the MAF density estimator. From these panels, we observe accurate recovery of seizure envelopes in pathological regions. Figure 6**E** and **F** show that the posterior distribution of heterogeneous η_i, and global coupling parameter G, respectively, indicating 100% accurate recovery of the true values (in green). Figure 6**G** confirms the reliability and accuracy of the estimates through shrinkage and z-score diagnostics. With our efficient implementation, generating 10k whole-brain simulations took less than a minute (using 10 CPU cores). The training took approximately 13 minutes to converge, while posterior sampling required only a few seconds. See Figure S6 for a similar analysis with a faster time separation ($\tau = 10\ ms$). These results demonstrate an ideal and fast Bayesian estimation, despite the stiffness of equations in each region and the high dimensionality of the parameters. See Figure S7 and Figure S8 showing the accuracy and reliability of estimation under different levels of additive and dynamical noise. Note that for the VEP model, the total integration time is less than 100 ms, and due to the model's stable behavior and a large time step integration, the simulation cost is significantly lower compared to other whole-brain models.

Targeting the fMRI data, we demonstrate the inference on the whole-brain dynamics using Montbrió model (see subsubsection 2.1.5). Figure 7 demonstrates the inference on heterogeneous control parameters of the Montbrió model, operating in a bistable regime (Table 5). Figure 7**A** and **B** show the observed and predicted BOLD time series, respectively, while Figure 7**C** and **D** illustrate the observed and predicted data features, such as the static and dynamical functional connectivity matrices (FC and FCD, respectively). Figure 7**E** and **F** show the estimated posterior distributions of excitability η_i per brain region, and the global coupling parameter G. Figure 7**G** displays the reliability and accuracy of estimation through the evaluation of posterior shrinkage and z-score (see Equation 13 and 14). See Figure S9 for estimation over different configurations of the ground-truth values in this model.

Due to the large number of simulations for training and the informativeness of the data features (both spatio-temporal and functional data features), the results indicate that we achieved accurate parameter estimation, consequently, a close agreement between the observed and predicted features of BOLD data. This required 500k simulations for training (see Figure S10), given the uniform priors $G \in \mathcal{U}(0, 1)$ and $\eta_i \in \mathcal{U}(-6, -3.5)$. After approximately 10 h of training (of MAF density estimator), posterior sampling took only 1 min. Our results indicate that training the MAF model was 2 to 4 times faster than the NSF model (see Figure S1**B**). Note that removing the spatio-temporal features and considering only FC/FCD as the data fea-

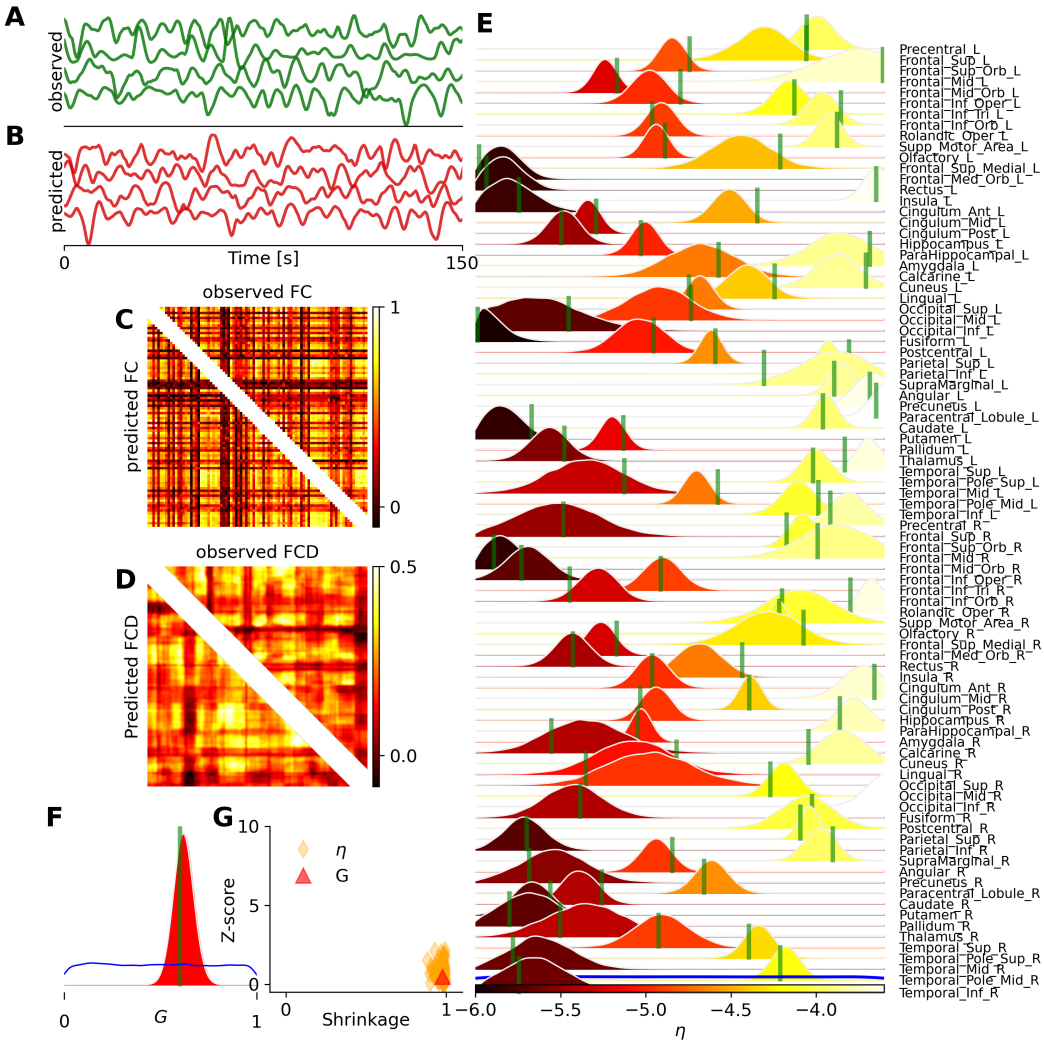

Figure 7. Bayesian inference on heterogeneous control parameters in the whole-brain dynamics using Montbrió model. The set of inferred parameters is $\vec{\theta} = \{G, \eta_i\} \in \mathbb{R}^{89}$, as the global scaling parameter and excitability per region, with $i \in \{1, 2, ..., N_n = 88\}$ parcelled regions. `VBI` provides accurate and reliable posterior estimation using both spatio-temporal and functional data features for training, with a budget of 500k simulations. (**A**) and (**B**) illustrate the observed and predicted BOLD signals, respectively. (**C**) and (**D**) show the observed (upper triangular) and predicted (lower triangular) data features (FC and FCD), respectively. (**E**) and (**F**) display the posterior distribution of excitability parameters η_i per region, and global coupling G, respectively. The true values and prior are shown as vertical green lines and a blue distribution, respectively. (**G**) shows the inference evaluation by the shrinkage and z-score of the posterior distributions.

tures (see Figure S11) leads to poor estimation of the excitability parameter across brain regions (see Figure S12). Interestingly, accurate estimation of the only global coupling parameter, $\vec{\theta} = \{G\} \in \mathbb{R}^1$, from only FC/FCD requires around 100 simulations (see Figure S13 and Figure S14). This showcase demonstrates the capability of `VBI` in inferring heterogeneous excitability, given the bistable brain dynamics, for fMRI studies. See Figure S15 and Figure S16 showing the accuracy and reliability of estimation under different levels of additive and dynamical noise.

Finally, in Figure 8, we show the inference on the so-called pDMF model, i.e., a whole-brain network model of the reduced Wong-Wang equation (see subsubsec-

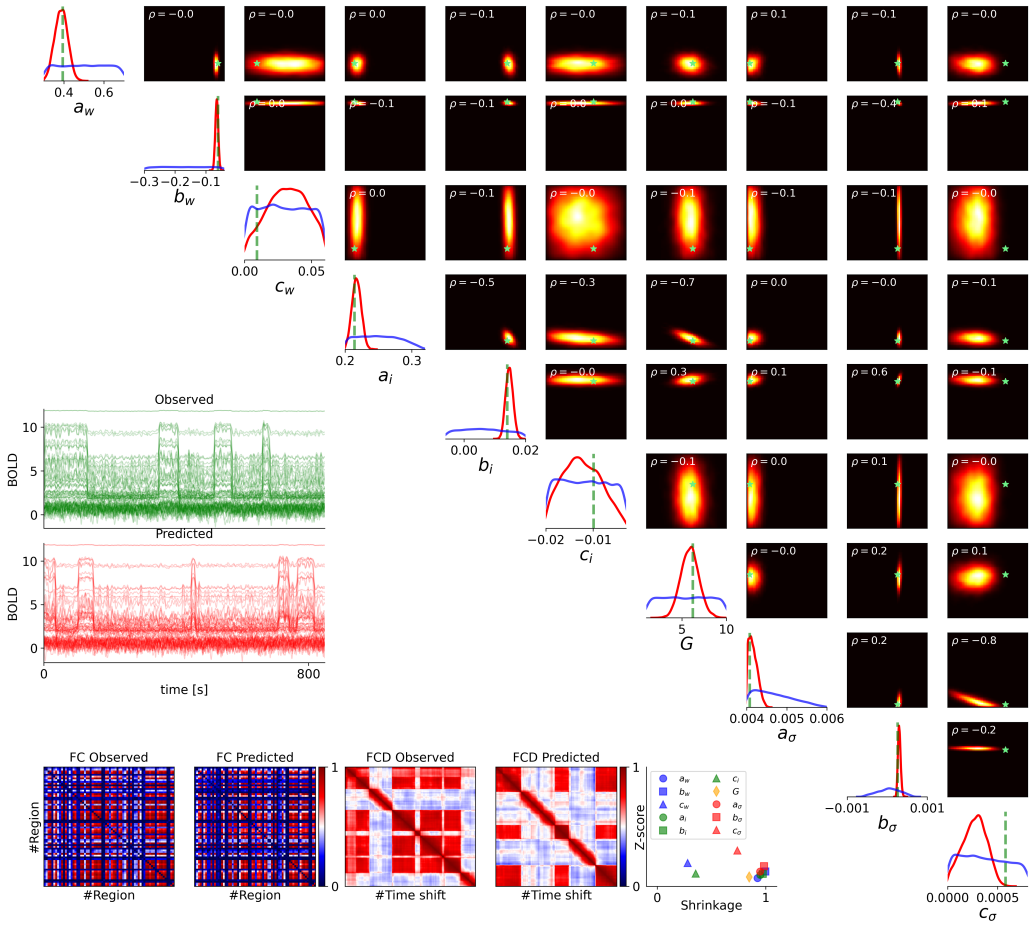

Figure 8. Bayesian inference on the parametric mean-field model of Wong-Wang (also known as pDMF model), with linear coefficients $(a_w, b_w, c_w, a_I, b_I, c_I, a_\sigma, b_\sigma, c_\sigma) \in \mathbb{R}^9$, reparameterizing the recurrent connection strength w_i, external input current I_i, and noise amplitude σ_i for each region. Summary statistics of spatio-temporal and functional data features were used for training, with a budget of 50k simulations. (**A**) The diagonal panels display the ground-true values (in green), the uniform prior (in blue), and the estimated posterior distributions (in red). The upper diagonal panels illustrate the joint posterior distributions between parameters, along with their correlation (ρ, in the upper left corners), and ground-truth values (green stars). High-probability areas are color-coded in yellow, while low-probability areas are shown in black. (**B**) The observed and predicted BOLD time series (in green and red, respectively). (**C**) The observed and predicted data features, such as FC/FCD matrices. (**D**) The inference evaluation by calculating the shrinkage and z-score of the estimated posterior distributions.

tion 2.1.6), comprising 10 control parameters: the global scaling of connections G and the linear coefficients $(a_w, b_w, c_w, a_I, b_I, c_I, a_\sigma, b_\sigma, c_\sigma) \in \mathbb{R}^9$. These parameters are introduced to reduce the dimension of whole-brain parameters as recurrent connection strength w_i, external input current I_i, and noise amplitude σ_i for each region (in total, 264 parameters were reduced to 9 dimensions; see Equation 7). Here, we used summary statistics of both spatio-temporal and functional data features extracted from simulated BOLD data to train the MAF density estimator, with a budget of 50k simulations. The training took around 160 min to converge, whereas posterior sampling took only a few seconds.

The diagonal panels in Figure 8**A** show estimated posterior distributions (in red), along with the prior (in blue) and true values (in green). The upper diagonal panels

illustrate the joint posterior distributions between parameters (i.e., statistical dependency between parameters). Figure 8**B** illustrates the observed and predicted BOLD time series, generated by true and estimated parameters (in blue and red, respectively). From Figure 8**C**, we can see a close agreement between the observed and predicted data features (FC/FCD matrices). Note that due to the stochastic nature of the generative process, we do not expect an exact element-wise correspondence between these features, but rather a match in their summary statistics, such as the mean, variance, and higher-order moments (see Figure S17). Figure 8**D** shows the posterior z-score versus shrinkage, indicating less accurate estimation for the coefficients c_w, c_I, and c_σ compared to others, as they are not informed by anatomical data such as the T1w/T2w myelin map and the first FC principal gradient (see Equation 7). This showcase demonstrates the advantage of Bayesian inference over optimization in assessing the accuracy and reliability of parameter estimation, whether informed by anatomical data.

Note that in the whole-brain network of Wong-Wang model (6), the global scaling parameter G and synaptic coupling J exhibit structural non-identifiability, meaning their combined effects on the system cannot be uniquely disentangled (see Figure S18, and Figure S19). This is evident in the parameter estimations corresponding to selected observations, where the posterior distributions appear diffuse. The joint posterior plots reveal a nonlinear dependency (banana shape) between G and J, arising from their product in the neural mass equation (see Equation 6). Such a nonlinear relationship between parameters poses challenges for deriving causal conclusions, as often occurs in other neural mass models. This is a demonstration of how Bayesian inference facilitates causal hypothesis testing without requiring additional non-identifiability analysis.

4. Discussion

This study introduces the Virtual Brain Inference (`VBI`), a flexible and integrative toolkit designed to facilitate probabilistic inference on complex whole-brain dynamics using connectome-based models (forward problem) and simulation-based inference (inverse problem). The toolkit leverages high-performance programming languages (C++) and dynamic compilers (such as Python's JIT compiler), alongside the computational power of parallel processors (GPUs), to significantly enhance the speed and efficiency of simulations. Additionally, `VBI` integrates popular feature extraction libraries with parallel multiprocessing to efficiently convert simulated time series into low-dimensional summary statistics. Moreover, `VBI` incorporates state-of-the-art deep neural density estimators (such as MAF and NSF generative models) to estimate the posterior density of control parameters within whole-brain models given low-dimensional data features. Our results demonstrated the versatility and efficacy of the `VBI` toolkit across commonly used whole-brain network models, such as Wilson-Cowan, Jansen-Rit, Stuart-Landau, Epileptor, Montbrió, and Wong-Wang equations placed at each region. The ability to perform parallel and rapid simulations, coupled with a taxonomy of feature extraction, allows for detailed and accurate parameter estimation from associated neuroimaging modalities such as (s)EEG/MEG/fMRI. This is crucial for advancing our understanding of brain dynamics and the underlying mechanisms of various brain disorders. Overall, `VBI` represents a substantial improvement over alternative methods, offering a robust framework for both simu-

lation and parameter estimation, and contributing to the advancement of network neuroscience, potentially, to precision medicine.

The alternatives for parameter estimation include optimization techniques [66, 34, 36, 63, 64], approximate Bayesian computation (ABC) method, and MCMC sampling. Optimization techniques are sensitive to the choice of the objective function (e.g., minimizing distance error or maximizing correlation) and do not provide estimates of uncertainty. Although multiple runs and thresholding can be used to address these issues, such methods often fall short in revealing relationships between parameters, such as identifying degeneracy, which is crucial for reliable causal inference. Alternatively, a technique known as ABC compares observed and simulated data using a distance measure based on summary statistics [147, 148, 149]. It is known that ABC methods suffer from the curse of dimensionality, and their performance also depends critically on the tolerance level in the accepted/rejected setting [71, 70]. The self-tuning variants of MCMC sampling have also been used for model inversion at the whole-brain level [76, 67]. Although MCMC is unbiased and exact with infinite runs, it can be computationally prohibitive, and sophisticated reparameterization methods are often required to facilitate convergence at whole-brain level [123, 150]. This becomes more challenging for gradient-based MCMC algorithms, due to the bistability and stiffness of neural mass models. Tailored to Bayes' rule, SBI sidesteps these issues by relying on expressive deep neural density estimators (such as MAF and NSF) on low-dimensional data features to efficiently approximate the posterior distribution of model parameters. Taking spiking neurons as generative models, this approach has demonstrated superior performance compared to alternative methods, as it does not require model or data features to be differentiable [71, 81].

In previous studies, we demonstrated the effectiveness of SBI on virtual brain models of neurological [68, 49, 50, 54], and neurodegenerative diseases [69, 51, 23] as well as focal intervention [59] and healthy aging [37]. In this work, we extended this probabilistic methodology to encompass a broader range of whole-brain network models, highlighting its flexibility and scalability in leveraging diverse computational resources, from CPUs/GPUs to high-performance computing facilities.

Our results indicated that the `VBI` toolkit effectively estimates posterior distributions of control parameters in whole-brain network modeling, offering a deeper understanding of the mechanisms underlying brain activity. For example, using the Montbrio and Wong-Wang models, we achieved a close match between observed and predicted FC/FCD matrices derived from BOLD time series (Figure 7 and Figure 8). Additionally, the Jansen-Rit and Stuart-Landau models provided accurate inferences of PSD from neural activity (Figure 4, and Figure 5), while the Epileptor model precisely captured the spread of of seizure envelopes (Figure 6). These results underscore the toolkit's capability to manage complex, high-dimensional data with precision. Uncertainty quantification using `VBI` can illuminate and combine the informativeness of data features (e.g., FC/FCD) and reveal the causal drivers behind interventions [37, 69, 59]. This adaptability ensures that `VBI` can be applied across various (source-localized) neuroimaging modalities, accommodating different computational capabilities and research needs.

Note that there is no specific rule for determining the optimal number of simulations required for training. In general, a larger number of simulations, depending on the available computational resources, tends to improve the quality of posterior estimation. However, when using synthetic data, we can monitor the z-score and pos-

terior shrinkage to assess the accuracy and reliability of the inferred parameters (see Figure S10). This also critically depends on the parameter dimensionality. For instance, in estimating only global coupling parameter, a maximum of 300 simulations was used, demonstrating accurate estimation across models and different configurations (see Figure S20), except for the Jansen-Rit model, where coupling did not induce a significant change in the intrinsic frequency of regional activity. Importantly, the choice of data features is critical, and some factors (e.g., that lead to inaccurate feature calculation) can lead to the collapse of this method. For instance, high noise levels in observations or dynamical noise can compromise the accurate calculation of data features, undermining the inference process (see Figure S7, Figure S8, Figure S15, Figure S16). Identifying the set of low-dimensional data features that are relevant to the control parameters for each case study is another challenge in effectively applying SBI. Nevertheless, the uncertainty of the posterior informs us about the predictive power of these features. Statistical moments of time series could be effective candidates for most models. However, this poses a formidable challenge for inference from empirical data, as certain moments, such as the mean and variance, may be lost during preprocessing steps. The hyperparameter and noise estimation can also be challenging for SBI. Moreover, there is no established rule for determining the number of simulations for training, aside from relying on z-score values during in-silico testing, as it depends on available computational resources.

Various sequential methods, such as SNPE [125], SNLE [151], and SNRE [152], have been proposed to reduce computational costs of SBI by iteratively refining the fit to specific targets. These approaches aim for more precise parameter estimation by progressively adjusting the model based on each new data set or subset, potentially enhancing the accuracy of the fit at the reduced computational effort. The choice of method depends on the specific characteristics and requirements of the problem being addressed [128]. Our previous study indicates that for inferring whole-brain dynamics of epilepsy spread, the SNPE method outperforms alternative approaches [68]. Nevertheless, sequential methods can become unstable, with simulators potentially diverging and causing probability mass to leak into regions that lack prior support [68]. In this study, we used single-round training to benefit from an amortization strategy. This approach brings the costs of simulation and network training upfront, enabling inference on new data to be performed rapidly (within seconds). This strategy facilitates personalized inference at the subject level, as the generative model is tailored by the SC matrix, thereby allowing for rapid hypothesis evaluation specific to each subject (e.g., in delineating the epileptogenic and propagation zones). Note that model comparison across different configurations or model structures, as well-established in dynamic causal modeling [153, 154, 155], has yet to be explored in this context.

Deep learning algorithms are increasingly gaining traction in the context of whole-brain modeling. The `VBI` toolkit leverages a class of deep generative models, called Normalizing Flows (NFs; [130, 131]), to model probability distributions given samples drawn from those distributions. Using NFs, a base probability distribution (e.g., a standard normal) is transformed into any complex distribution (potentially multimodal) through a sequence of invertible transformations. Variational autoencoders (VAEs; [156, 157]) is a class of deep generative models to encode data into a latent space and then decode it back to reconstruct the original data. Recently, Sip et al. [158] introduced a method using VAEs for nonlinear dynamical system identifica-

tion, enabling the inference of neural mass models and region- and subject-specific parameters from functional data. VAEs have also been employed for dimensionality reduction of whole-brain functional connectivity [42], and to investigate various pathologies and their severity by analyzing the evolution of trajectories within a low-dimensional latent space [52]. Additionally, Generative Adversarial Networks (GANs; [159, 160]) have demonstrated remarkable success in mapping latent space to data space by learning a manifold induced from a base density [129]. This method merits further exploration within the context of whole-brain dynamics. To fully harness the potential of deep generative models in large-scale brain network modeling, integrating VAEs and GANs into the `VBI` framework would be beneficial. This will elucidate their strengths and limitations within this context and guide future advancements in the field.

In summary, `VBI` offers fast simulations, taxonomy of feature extraction, and deep generative models, making it a versatile tool for model-based inference from different neuroimaging modalities, helping researchers to explore brain (dys)functioning in greater depth. This advancement not only enhances our theoretical understanding but also holds promise for practical applications in diagnosing and treating neurological conditions.

5. Glossary of technical terms

Bayesian Rule: A fundamental belief updating principle that calculates the probability of a hypothesis given new evidence.

Deep neural density estimators: A class of artificial neural network-based approaches that are used to learn and approximate the underlying probability distribution from a given dataset.

Control (generative) parameters: The bifurcation parameters, setting, or configuration within a generative model that controls the synthesis of data and potentially represents causal relationships.

Generative model: A statistical, machine learning, or mechanistic model that represents the underlying data distribution to generate new data resembling the original dataset.

Likelihood: The conditional probability of observing the evidence given a particular hypothesis.

Markov chain Monte Carlo: A family of stochastic algorithms used for uncertainty quantification by drawing random samples from probability distributions, in which the sampling process does not require knowledge of the entire distribution.

Prior: The initial probability assigned to a hypothesis before considering new evidence.

Probability distribution: The statistical description of potential outcomes of random events, where a numerical measure is assigned to the possibility of each specific outcome.

Probability density estimation: The process of inferring the underlying probability distribution of a random event based on observed data.

Posterior: The updated probability of a hypothesis after taking into account both prior beliefs and observed evidence.

Simulation-based inference: A machine learning approach that involves generating synthetic data through forward simulations to make inferences about complex systems, often when analytic or computational solutions are unavailable.

Whole-brain or virtual brain models: Computational models informed by personalized anatomical data, i.e., a set of equations describing regional brain dynamics placed at each node, which are then connected through structural connectivity matrix.

Information Sharing Statement

All code is available on GitHub (https://github.com/ins-amu/vbi).

Acknowledgements

This project/research has received funding from the European Union's Horizon Europe Programme under the Specific Grant Agreement No. 101147319 (EBRAINS 2.0 Project), No. 101137289 (Virtual Brain Twin Project), No. 101057429 (project environMENTAL), and government grant managed by the Agence Nationale de la Recherche reference ANR-22-PESN-0012 (France 2030 program). We acknowledge the use of Fenix Infrastructure resources, which are partially funded from the European Union's Horizon 2020 research and innovation programme through the ICEI project under the grant agreement No. 800858. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author contributions

Conceptualization: V.J., and M.H. Methodology: A.Z., M.W., L.D., and M.H. Software: A.Z, M.W., L.D., and M.H. Investigation: A.Z. Visualization: A.Z., Supervision: V.J., S.P., and M.H. Funding acquisition: V.J., M.H. Writing - original draft: A.Z., and M.H. Writing - review & editing: A.Z, M.W., L.D., S.P., V.J, and M.H.

6. Supplementary

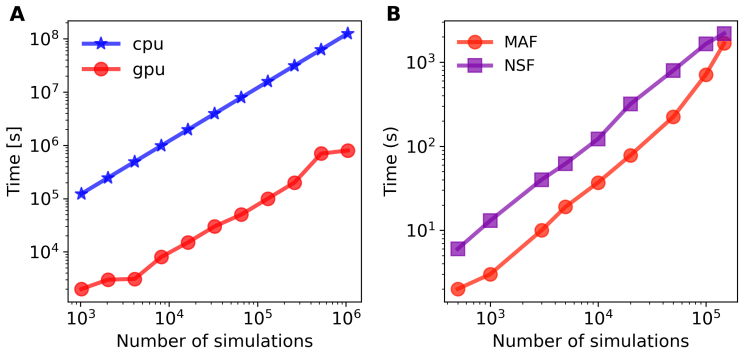

Figure S1. (**A**) Benchmarking the simulation cost of the whole-brain network of Montbrió model, with 2.5M time step on Intel(R) Core(TM) i9-10900 CPU 2.80 GHz (in red) and NVIDIA RTX A5000 GPU (in blue). GPUs deliver substantial speedups up to 100x over multi-core CPUs. (**B**) Comparing MAF and NSF density estimators, MAF is typically 2-4 times faster than NSF during the training process.

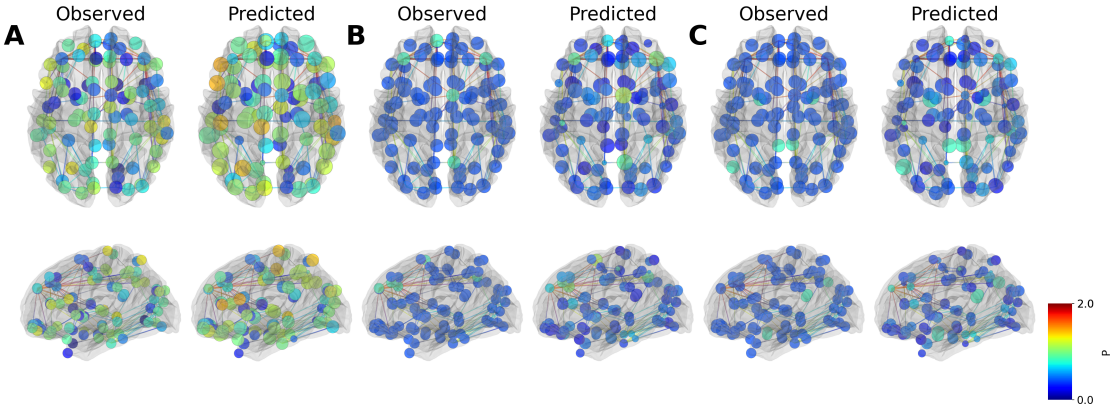

Figure S2. Inference across different configurations of heterogeneous generative parameters in the whole-brain network model based on the Wilson-Cowan framework. The set of unknown control parameters is $\vec{\theta} = \{G, P_i\} \in \mathbb{R}^{89}$. We observe accurate posterior estimation across different scenarios by varying the external input current map across brain regions. Features derived from the power spectrum density were used for training, with a simulation budget of 260k. The regions are color-coded based on their current values: blue indicates low, green medium, and red high external input current and the size of the regions reflects the shrinkage of the posterior distribution for the corresponding inferred parameters. Top: axial view. Bottom: sagittal view. The three configurations (from left to right) correspond to: randomly assigned heterogeneous P_i values, increased P_i for the five strongest nodes, and increased P_i for the five weakest nodes.

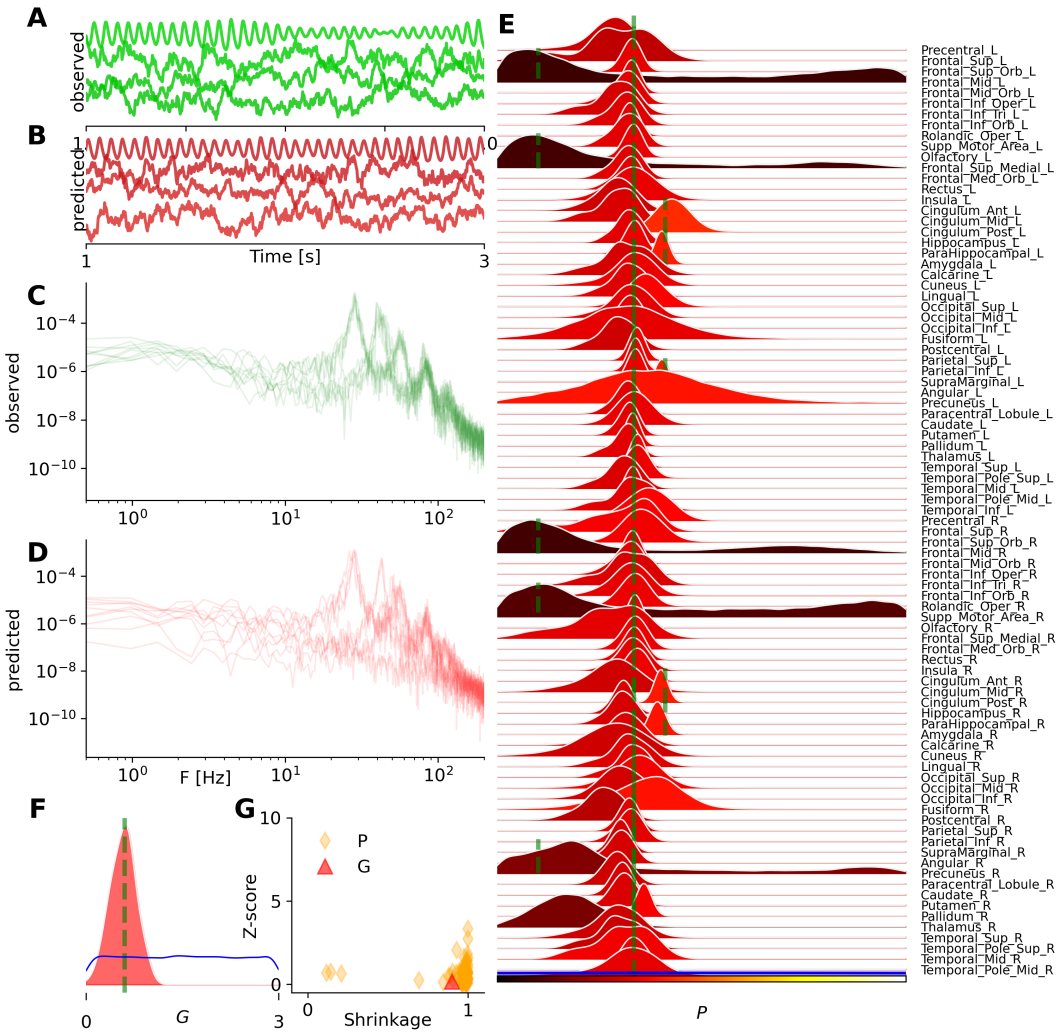

Figure S3. Bayesian inference on heterogeneous control parameters in the whole-brain network of Wilson-Cowan model, as Figure 3, but ignoring the spatial information of the data features. Here, we observe that the MAF density estimator leads to a multimodal posterior distribution for low values of P_i, when trained with a reduced set of data features, suggesting increased uncertainty or potential non-identifiability under limited information.

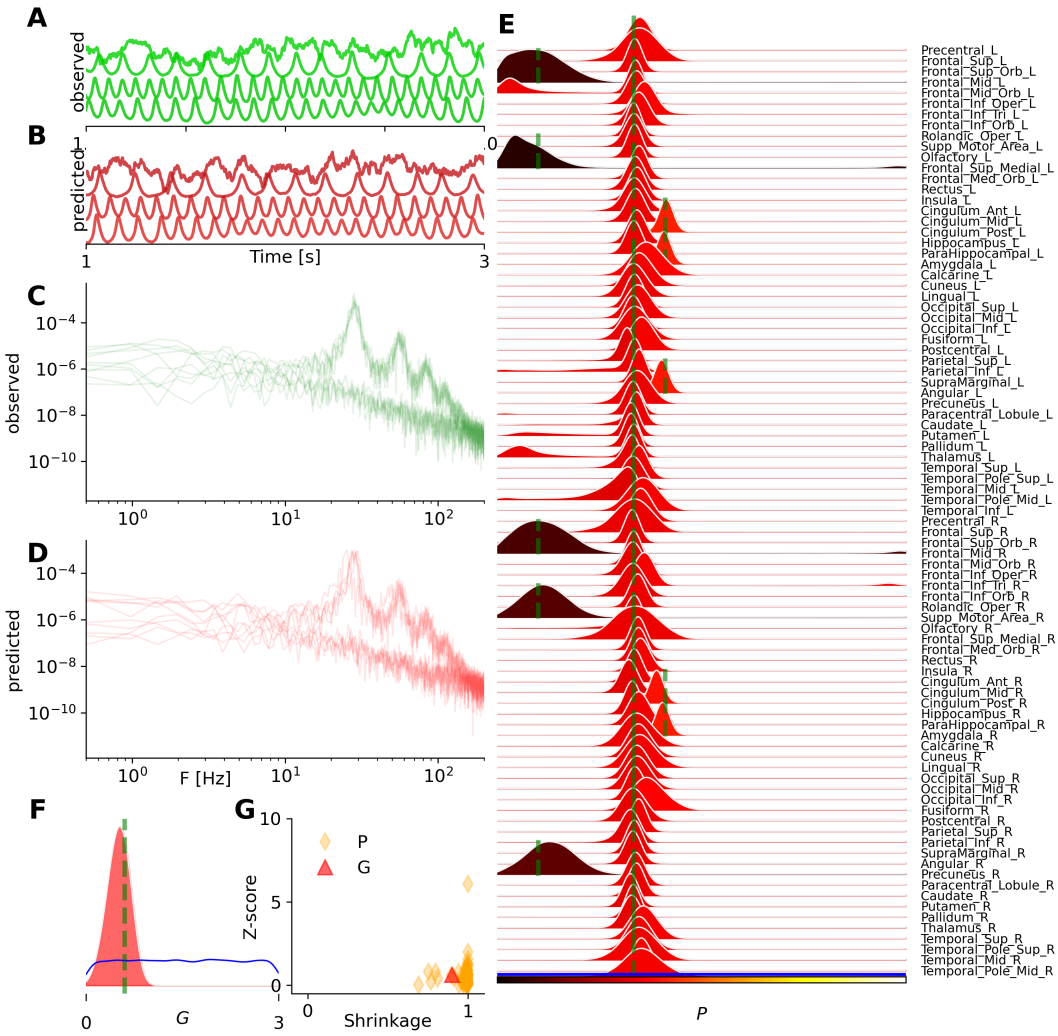

Figure S4. Same as Figure S3, but using the NSF density estimator. Here, we observe more accurate estimation, however, NSF required approximately 6 hours of training compared to 20 minutes for MAF.

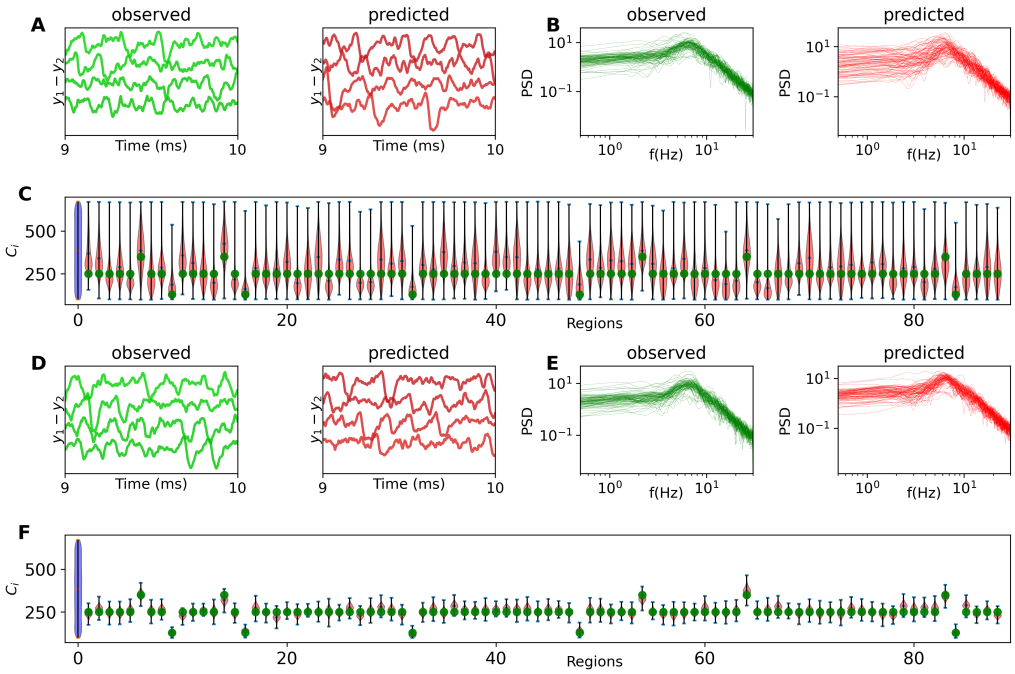

Figure S5. Inference on heterogeneous generative parameters in the whole-brain network of the Jansen-Rit model (Equation 2). The set of inferred parameters is $\vec{\theta} = \{G, C_i\} \in \mathbb{R}^{89}$, consisting of the global scaling parameter G and the average number of synapses between neural populations C_i per region, with $i \in \{1, 2, ..., N_n = 88\}$ parcelled regions. For panels **A** to **C**, only the area under the curve of the power spectral density (PSD) in the [0,30] Hz band was used, while for panels **D** to **F**, additional statistical features of the PSD were included in the training, with a total budget of 50k simulations. (**A**) and (**D**) illustrate the observed and predicted EEG signals. (**B**) and (**E**) show the observed and predicted PSD signals. (**C**) and (**F**) display the posterior distribution of C_i for each region. The true values are indicated by vertical green lines, and the prior distribution is shown in blue.

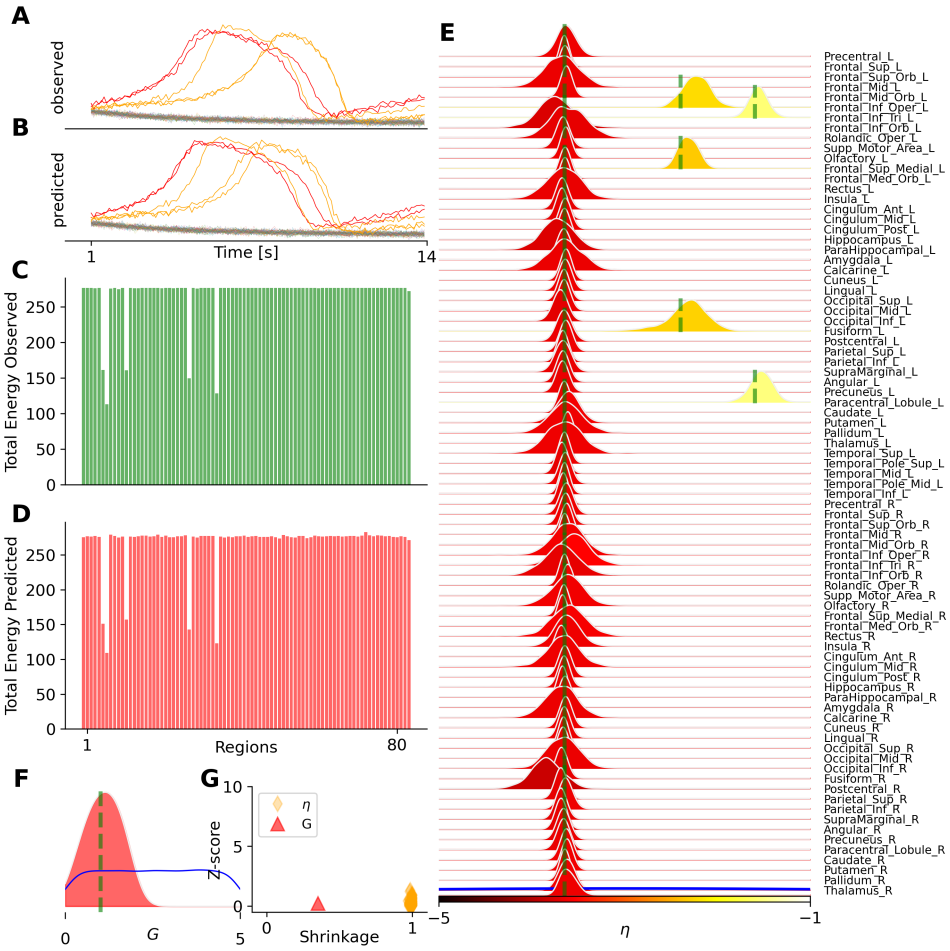

Figure S6. Bayesian inference on the spatial map of epileptogenicity across brain regions in the VEP model (Equation 4), with slow time scale separation, $\tau = 10\ ms$. (**A**) The observed seizure envelope generated by the 2D Epileptor model given two regions as epileptogenic zones (in red) and three regions as propagation zones (in yellow), while the rest are healthy (in green). (**B**) The predicted seizure envelope, by training MAF density estimator on a dataset containing 10k simulations, using only the total power per region as the data features. (**C**) and (**D**) show the observed and predicted data features, respectively. (**E**) and (**F**) show the posterior distribution of heterogeneous η_i, and global coupling parameter G, respectively. (**F**) The posterior z-scores versus posterior shrinkages for estimated parameters.

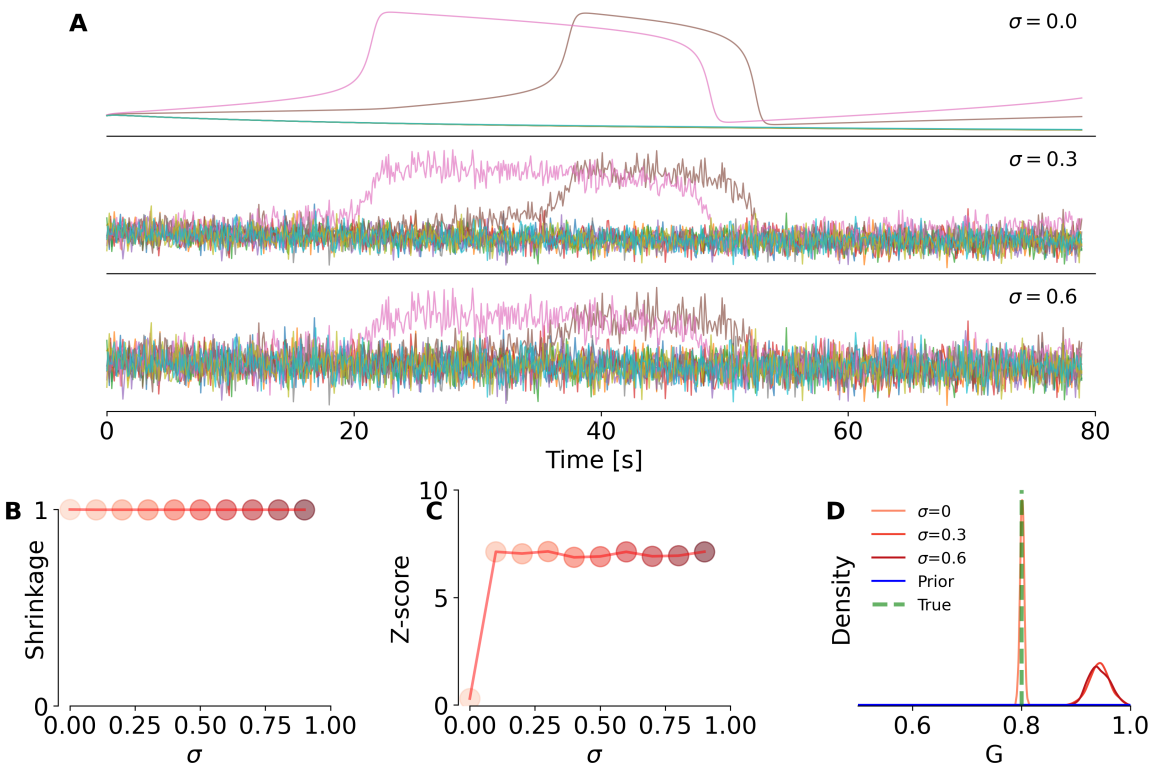

Figure S7. Evaluating the performance of estimation over the global coupling parameter, under different levels of additive noise in the VEP model. (**A**) Simulated time series under different levels of additive noise. (**B**) and (**C**) Posterior shrinkage and z-score, respectively, as function of noise level σ. (**D**) Ground-true (in green), prior (in blue), and posterior distributions (in red) for global coupling parameter. We observe that at higher noise levels (e.g., $\sigma = 0.3$ and $\sigma = 0.6$), the ground truth is not captured within the posterior, leading to increased z-scores compared to the noise-free case $\sigma = 0$.

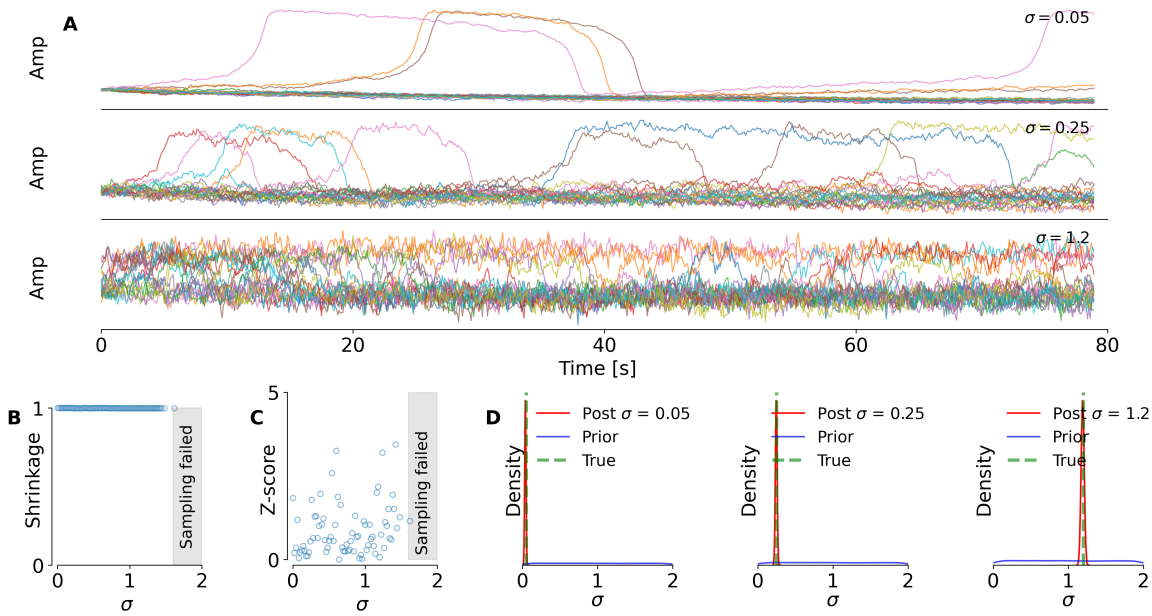

Figure S8. Evaluating the accuracy of dynamical noise estimation in the VEP model. (**A**) Simulated time series generated under different levels of dynamical noise. (**B**) and (**C**) Posterior shrinkage and z-score, respectively, across different scenarios, demonstrating the accuracy and reliability of estimating the dynamical noise parameter. (**D**) Ground-true (in green), prior (in blue) and posterior distributions (in red), for different realizations of the dynamical noise parameter σ in the system. We observe robust estimation even at high levels of dynamical noise, though for $\sigma \geq 1.6$ the posterior sampling fails (shaded area).

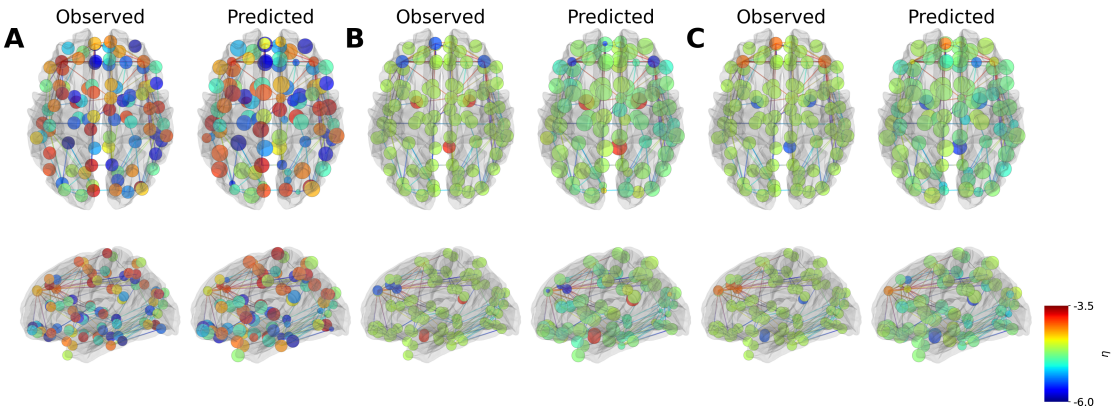

Figure S9. Inference over different configurations of heterogeneous generative parameters in the whole-brain network model of the Montbrió, (see Figure 7). The set of inferred parameters is $\vec{\theta} = \{G, \eta_i\} \in \mathbb{R}^{89}$, as the global scaling parameter and excitability per region, with $i \in \{1, 2, ..., N_n = 88\}$ parcelled regions.

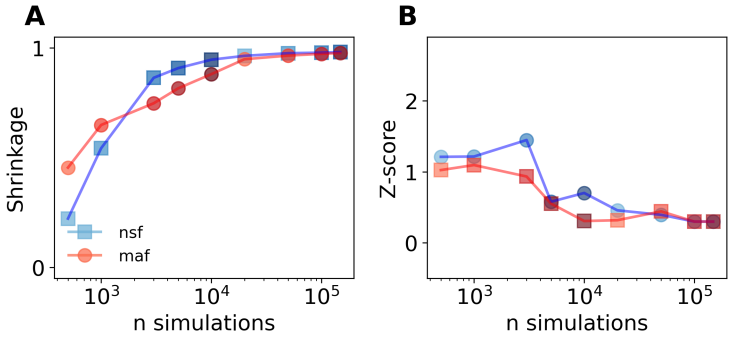

Figure S10. Evaluating the reliability and accuracy of estimation (parameter G) with respect to the number of simulation used for training the MAF/NSF density estimator by (**A**) posterior shrinkage, and (**B**) posterior z-score, versus the number of simulations.

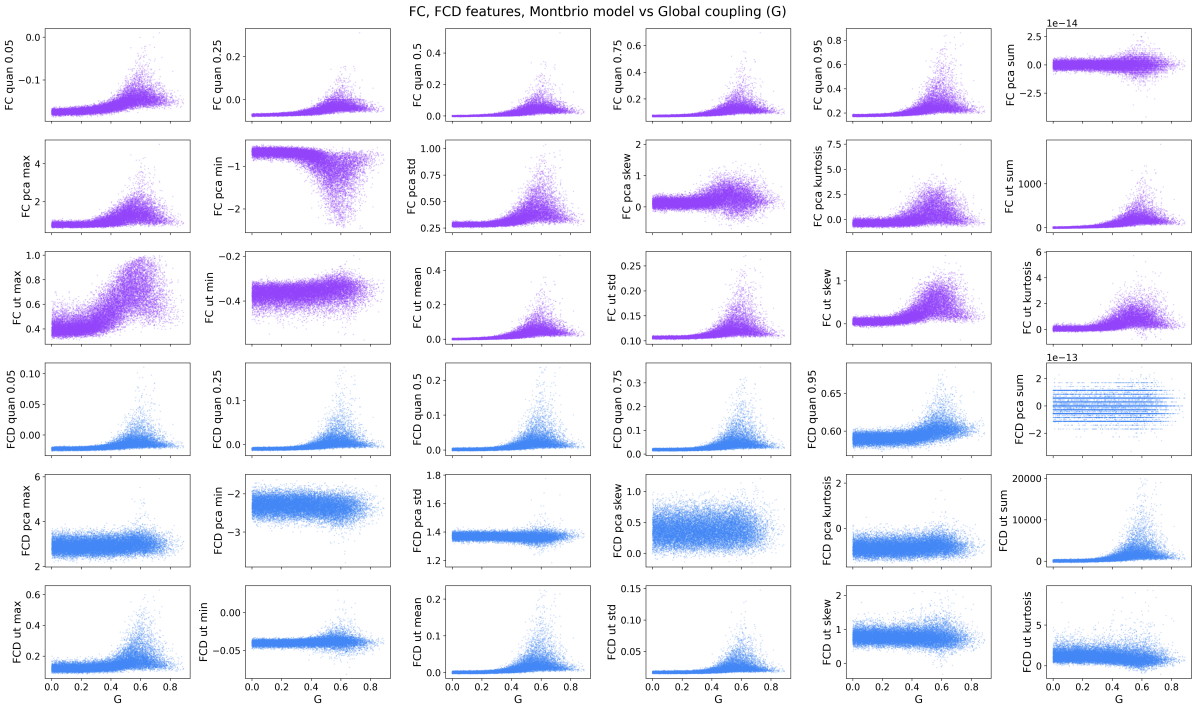

Figure S11. Scatter plot of summary statistics from the FC/FCD matrices in the whole-brain network of Montbrió model (Equation 5).

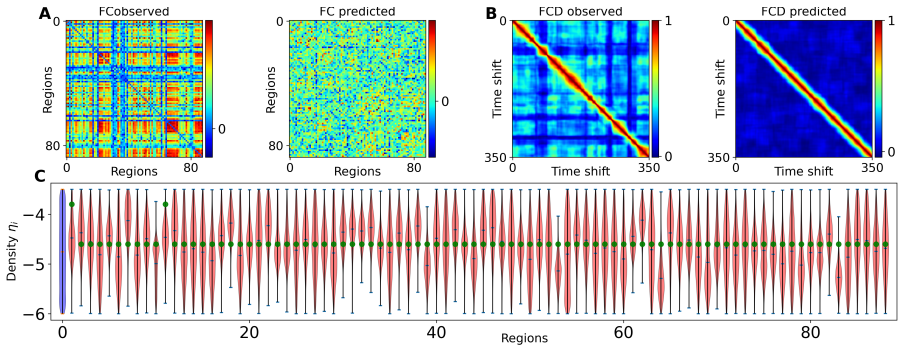

Figure S12. Inference on heterogeneous generative parameters in the whole-brain network of Montbrió model (Equation 5) involves estimating a parameter set denoted as $\vec{\theta} = \{G, \eta_i\} \in \mathbb{R}^{89}$ from FC/FCD matrices. Here, G is the global scaling parameter, and η_i represents the excitability of each brain region, where $i \in \{1, 2, \ldots, N_n = 88\}$ corresponds to the parcelled regions. Relying on only functional data features (FC/FCD) for training proves insufficient to accurately estimate the posterior distribution of heterogeneous excitabilities, even with a budget of 500k simulations. In (**A**), the posterior estimates are shown in red, the true values of η_i in green, and the prior distribution in blue. Panels (**B**) and (**C**) show the observed and predicted FC/FCD matrices, respectively, as the data features. The simulation to generate training data took approximately 24 hours, utilizing 10 GPUs concurrently. The training process lasted around 8 hours, while sampling required just a few seconds.

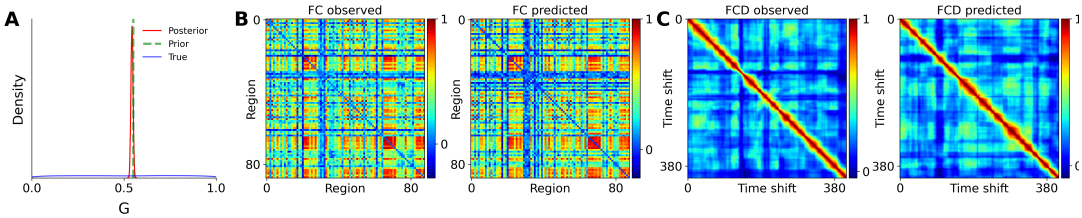

Figure S13. Inference of the global coupling parameter $\vec{\theta} = \{G\} \in \mathbb{R}^1$ in the whole-brain network of Montbrió model (Equation 5) using functional data features, using a budget of only 100 simulations. (**A**) The uniform prior distribution of G in blue, the estimated posterior in red and the true value is shown in green. (**B**) and (**C**) The observed and predicted FC/FCD matrices, respectively.

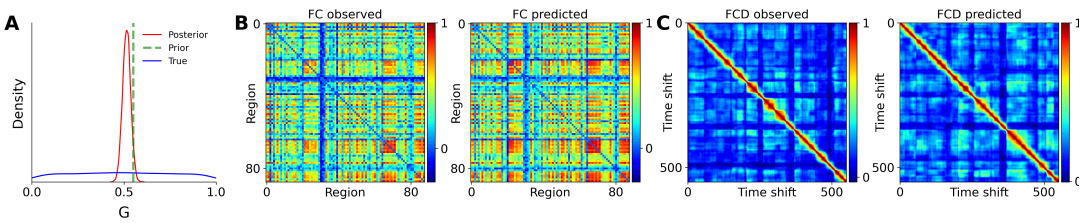

Figure S14. Inference of the global coupling parameter $\vec{\theta} = \{G\} \in \mathbb{R}^1$ in the whole-brain model of Montbrió (Equation 5) using TVB simulator, using summation of FC and FCD matrix elements as low dimensional data features, constrained by a budget of only 100 simulations. (**A**) The uniform prior distribution of G in blue, the estimated posterior in red and the true value is shown in green. (**B**) and (**C**) show the observed and predicted FC/FCD matrices, respectively.

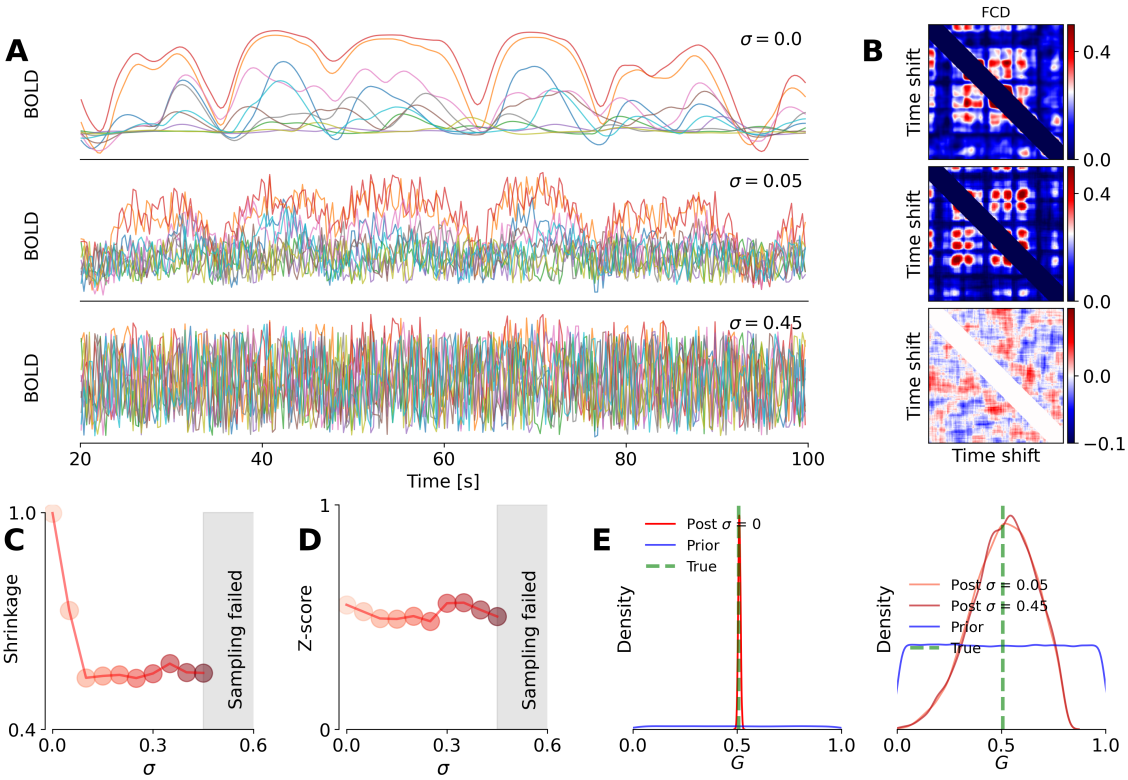

Figure S15. Evaluating the performance of estimation over the global coupling parameter under different levels of noise added to the BOLD signal in the Montbrió model. (**A**) and (**B**) Simulated time series and corresponding FCD matrices, respectively, under different levels of additive noise. (**C**) and (**D**) Posterior shrinkages and z-scores, respectively, for different noise values. (**E**) Ground-true (in green), prior (in blue) and posterior distribution (in red), for estimation of coupling parameter G, given different values of additive noise to the BOLD signals, indicating robust inference. However, for very high values of noise ($\sigma \geq 0.45$) the posterior sampling fails (shaded area).

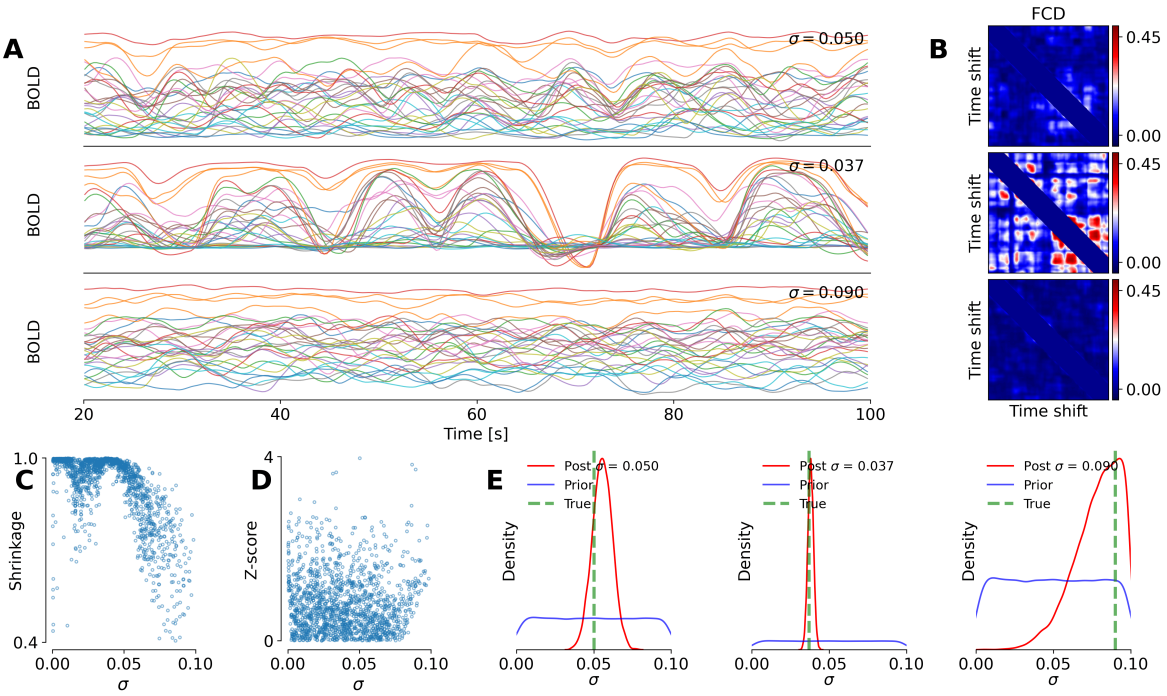

Figure S16. Evaluating the accuracy of dynamical noise estimation in the Montbrió model. (**A**) and (**B**) Simulated time series and corresponding FCD matrices, respectively, under different levels of dynamical noise. (**C**) and (**D**) Posterior shrinkages and z-scores, respectively, for different noise values. (**E**) Ground-true (in green), prior (in blue) and posterior distributions (in red), for different realizations of the dynamical noise parameter σ in the system. We observe robust estimation even at high levels of dynamical noise.

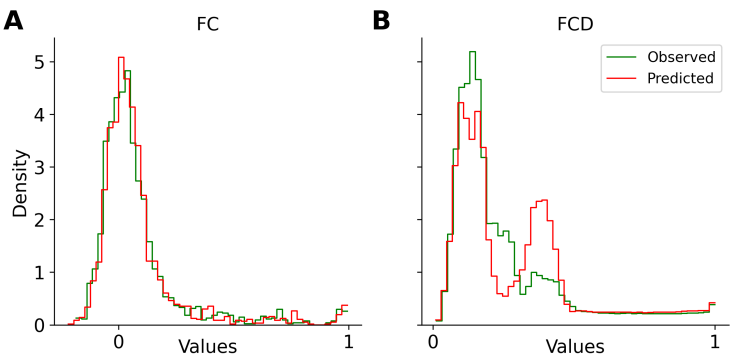

Figure S17. The distribution of observed and predicted functional connectivity (FC) in panel and the functional connectivity dynamics in panel (**B**), which are demonstrated in Figure 8.

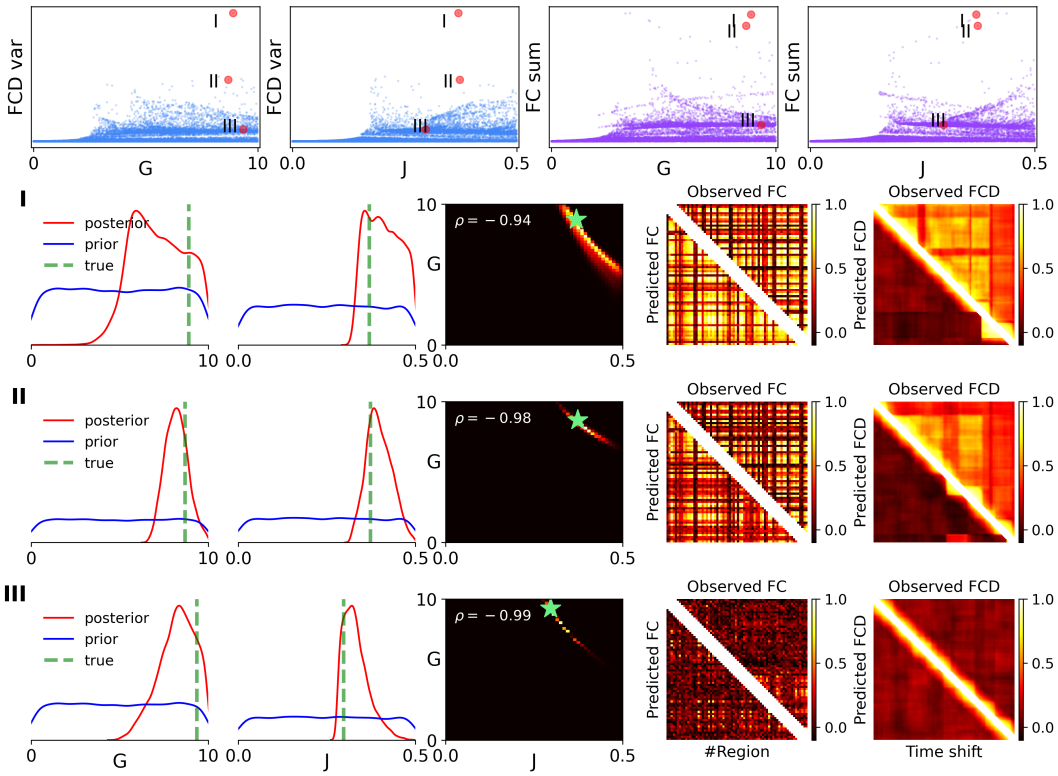

Figure S18. The non-identifiability in estimating G and J in the Wong-Wang model (Equation 6). A total of 16k simulations and functional data features (see Figure S18) were used for training. The upper row illustrates the FCD variance and the sum of FC matrices plotted against the parameters G and J, respectively, for three selected observations marked by red circles. The lower panels display the parameter estimations corresponding to the selected observations. Due to structural non-identifiability in the product of G and J, the posterior distributions appear diffuse, showing a nonlinear dependency (banana shape) in the joint posterior plots.

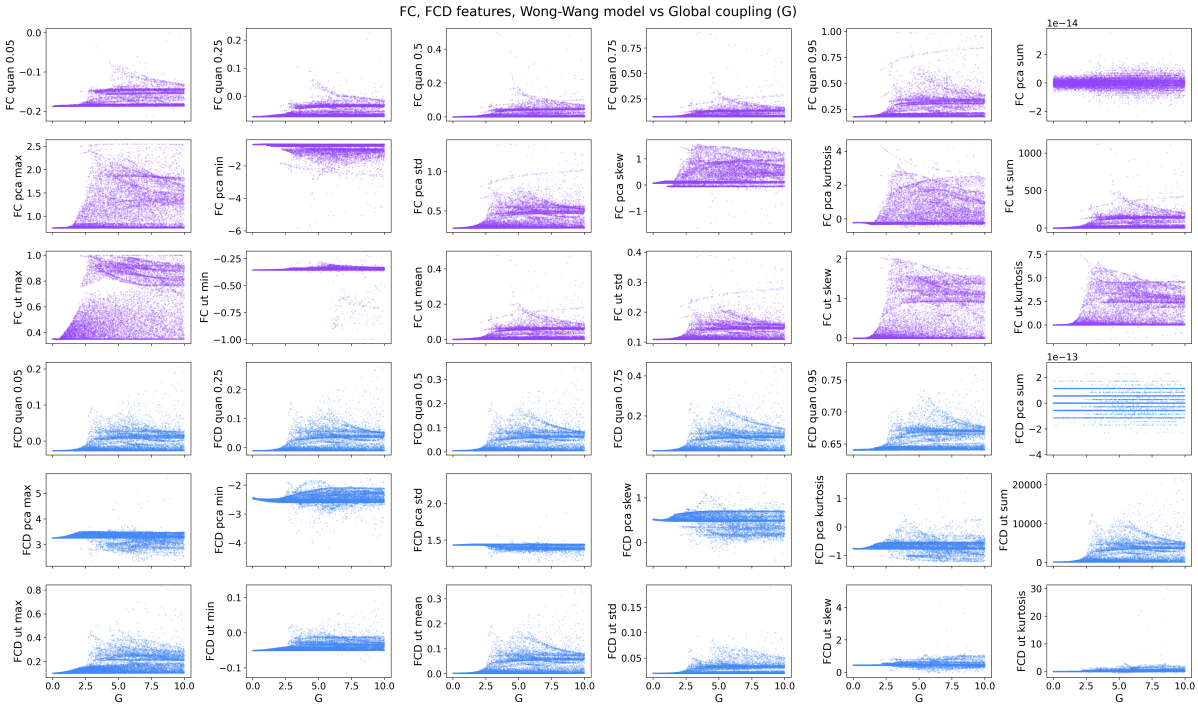

Figure S19. Scatter plot of summary statistics from the FC/FCD matrices in the whole-brain network of Wong-Wang model (Equation 6).

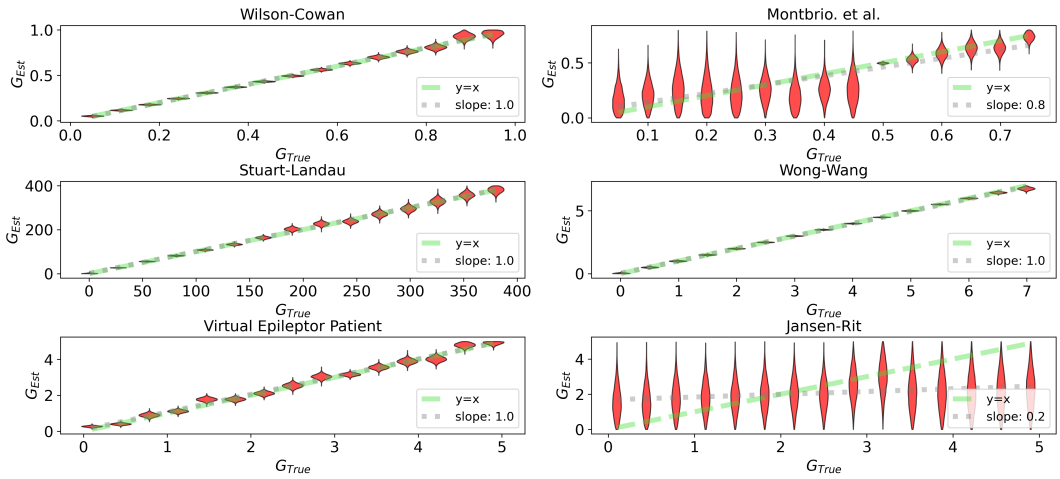

Figure S20. Estimation of the global coupling parameter across whole-brain models, evaluated for multiple configurations (each represented by a violin plot). The MAF density estimator was trained using the same features as described for the heterogeneous case in the main text. The simulation budget was limited to a maximum of 300. The green dashed line indicates an ideal estimation, while the gray dashed line represents a linear regression based on the mean values of the posterior sampling distributions. We can observe an accurate posterior estimation for coupling parameter across different models and configurations, except for the Jansen-Rit model, where coupling did not induce a significant change in the intrinsic frequency of regional activity. Note that the Montbrió model shows more diffuse estimations compared to the others, due to switching dynamics driven by noise.

References

[1] Maria I Falcon, Viktor Jirsa, and Ana Solodkin. A new neuroinformatics approach to personalized medicine in neurology: The virtual brain. *Current opinion in neurology*, 29(4):429–436, 2016.

[2] Lin Tan, Teng Jiang, Lan Tan, and Jin-Tai Yu. Toward precision medicine in neurological diseases. *Annals of translational medicine*, 4(6), 2016.

[3] Jacob W Vogel, Nick Corriveau-Lecavalier, Nicolai Franzmeier, Joana B Pereira, Jesse A Brown, Anne Maass, Hugo Botha, William W Seeley, Dani S Bassett, David T Jones, et al. Connectome-based modelling of neurodegenerative diseases: towards precision medicine and mechanistic insight. *Nature Reviews Neuroscience*, 24(10):620–639, 2023.

[4] Leanne M Williams and Susan Whitfield Gabrieli. Neuroimaging for precision medicine in psychiatry. *Neuropsychopharmacology*, pages 1–12, 2024.

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

[6] Huifang E Wang, Marmaduke Woodman, Paul Triebkorn, Jean-Didier Lemarechal, Jayant Jha, Borana Dollomaja, Anirudh Nihalani Vattikonda, Viktor Sip, Samuel Medina Villalon, Meysam Hashemi, et al. Delineating epileptogenic networks using brain imaging data and personalized modeling in drug-resistant epilepsy. *Science Translational Medicine*, 15(680):eabp8982, 2023.

[7] Lauren N Ross and Dani S Bassett. Causation in neuroscience: Keeping mechanism meaningful. *Nature Reviews Neuroscience*, 25(2):81–90, 2024.

[8] H.R. Wilson and J.D. Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.*, 12:1–24, 1972.

[9] V.K. Jirsa and H. Haken. Field theory of electromagnetic brain activity. *Phys. Rev. Lett.*, 77(5):960–963, 1996.

[10] G. Deco, V.K. Jirsa, P.A. Robinson, M. Breakspear, and K. Friston. The dynamic brain: from spiking neurons to neural masses and cortical fields. *PloS Comp. Biol.*, 4(8), 2008.

[11] Viktor K. Jirsa, William C. Stacey, Pascale P. Quilichini, Anton I. Ivanov, and Christophe Bernard. On the nature of seizure dynamics. *Brain*, 137(8):2210–2230, 2014.

[12] Ernest Montbrió, Diego Pazó, and Alex Roxin. Macroscopic description for networks of spiking neurons. *Physical Review X*, 5(2):021028, 2015.

[13] Blake J Cook, Andre DH Peterson, Wessel Woldman, and John R Terry. Neural field models: A mathematical overview and unifying framework. *Mathematical Neuroscience and Applications*, 2, 2022.

[14] Christopher J Honey, Olaf Sporns, Leila Cammoun, Xavier Gigandet, Jean-Philippe Thiran, Reto Meuli, and Patric Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009.

[15] Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: a structural description of the human brain. *PLoS computational biology*, 1(4):e42, 2005.

[16] Michael Schirner, Simon Rothmeier, Viktor K Jirsa, Anthony Randal McIntosh, and Petra Ritter. An automated pipeline for constructing personalized virtual brains from multimodal neuroimaging data. *NeuroImage*, 117:343–357, 2015.

[17] Vincent Bazinet, Justine Y Hansen, and Bratislav Misic. Towards a biologically annotated brain connectome. *Nature reviews neuroscience*, 24(12):747–760, 2023.

[18] 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.

[19] Michael Schirner, Lia Domide, Dionysios Perdikis, Paul Triebkorn, Leon Stefanovski, Roopa Pai, Paula Prodan, Bogdan Valean, Jessica Palmer, Chloê Langford, André Blickensdörfer, Michiel van der Vlag, Sandra Diaz-Pier, Alexander Peyser, Wouter Klijn, Dirk Pleiter, Anne Nahm, Oliver Schmid, Marmaduke Woodman, Lyuba Zehl, Jan Fousek, Spase Petkoski, Lionel Kusch, Meysam Hashemi, Daniele Marinazzo, Jean-François Mangin, Agnes Flöel, Simisola Akintoye, Bernd Carsten Stahl, Michael Cepic, Emily Johnson, Gustavo Deco, Anthony R. McIntosh, Claus C. Hilgetag, Marc Morgan, Bernd Schuller, Alex Upton, Colin McMurtrie, Timo Dickscheid, Jan G. Bjaalie, Katrin Amunts, Jochen Mersmann, Viktor Jirsa, and Petra Ritter. Brain simulation as a cloud service: The virtual brain on ebrains. *NeuroImage*, 251:118973, 2022.

[20] Katrin Amunts, Javier DeFelipe, Cyriel Pennartz, Alain Destexhe, Michele Migliore, Philippe Ryvlin, Steve Furber, Alois Knoll, Lise Bitsch, Jan G Bjaalie, et al. Linking brain structure, activity, and cognitive function through computation. *Eneuro*, 9(2), 2022.

[21] Egidio D'Angelo and Viktor Jirsa. The quest for multiscale brain modeling. *Trends in neurosciences*, 45(10):777–790, 2022.

[22] Gustavo Patow, Ignacio Martin, Yonatan Sanz Perl, Morten L Kringelbach, and Gustavo Deco. Whole-brain modelling: an essential tool for understanding brain dynamics. *Nature Reviews Methods Primers*, 4(1):53, 2024.

[23] Meysam Hashemi, Damien Depannemaecker, Marisa Saggio, Paul Triebkorn, Giovanni Rabuffo, Jan Fousek, Abolfazl Ziaeemehr, Viktor Sip, Anastasios Athanasiadis, Martin Breyton, Marmaduke Woodman, Huifang Wang, Spase Petkoski, Pierpaolo Sorrentino, and Viktor Jirsa. Principles and operation of virtual brain twins. *IEEE Reviews in Biomedical Engineering*, 14:1–29, 2025.

[24] Olaf Sporns. *Networks of the Brain.* MIT press, 2016.

[25] Danielle S Bassett and Olaf Sporns. Network neuroscience. *Nature neuroscience*, 20(3):353–364, 2017.

[26] Anandamohan Ghosh, Y Rho, Anthony Randal McIntosh, Rolf Kötter, and Viktor K Jirsa. Noise during rest enables the exploration of the brain's dynamic repertoire. *PLoS computational biology*, 4(10):e1000196, 2008.

[27] Christopher J Honey, Jean-Philippe Thivierge, and Olaf Sporns. Can structure predict function in the human brain? *Neuroimage*, 52(3):766–776, 2010.

[28] Hae-Jeong Park and Karl Friston. Structural and functional brain networks: from connections to cognition. *Science*, 342(6158):1238411, 2013.

[29] Francesca Melozzi, Eyal Bergmann, Julie A. Harris, Itamar Kahn, Viktor Jirsa, and Christophe Bernard. Individual structural features constrain the mouse functional connectome. *Proceedings of the National Academy of Sciences*, 116(52):26961–26969, 2019.

[30] 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.

[31] Guozheng Feng, Yiwen Wang, Weijie Huang, Haojie Chen, Jian Cheng, and Ni Shu. Spatial and temporal pattern of structure–function coupling of human brain connectome with development. *eLife*, 13:RP93325, 2024.

[32] Jacob Tanner, Joshua Faskowitz, Andreia Sofia Teixeira, Caio Seguin, Ludovico Coletta, Alessandro Gozzi, Bratislav Mišić, and Richard F Betzel. A multimodal, asymmetric, weighted, and signed description of anatomical connectivity. *Nature Communications*, 15(1):5865, 2024.

[33] G. Deco, V.K. Jirsa, and A.R. McIntosh. Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience*, 12(1):43–56, 2011.

[34] Peng Wang, Ru Kong, Xiaolu Kong, Raphaël Liégeois, Csaba Orban, Gustavo Deco, Martijn P Van Den Heuvel, and BT Thomas Yeo. Inversion of a large-scale circuit model reveals a cortical hierarchy in the dynamic resting human brain. *Science advances*, 5(1):eaat7854, 2019.

[35] Abolfazl Ziaeemehr, Mina Zarei, Alireza Valizadeh, and Claudio R Mirasso. Frequency-dependent organization of the brain's functional network through delayed-interactions. *Neural Networks*, 132:155–165, 2020.

[36] Xiaolu Kong, Ru Kong, Csaba Orban, Peng Wang, Shaoshi Zhang, Kevin Anderson, Avram Holmes, John D Murray, Gustavo Deco, Martijn van den Heuvel, et al. Sensory-motor cortices shape functional connectivity dynamics in the human brain. *Nature communications*, 12(1):1–15, 2021.

[37] Mario Lavanga, Johanna Stumme, Bahar Hazal Yalcinkaya, Jan Fousek, Christiane Jockwitz, Hiba Sheheitli, Nora Bittner, Meysam Hashemi, Spase Petkoski, Svenja Caspers, et al. The virtual aging brain: Causal inference supports interhemispheric dedifferentiation in healthy aging. *NeuroImage*, 283:120403, 2023.

[38] Shaoshi Zhang, Bart Larsen, Valerie J Sydnor, Tianchu Zeng, Lijun An, Xiaoxuan Yan, Ru Kong, Xiaolu Kong, Ruben C Gur, Raquel E Gur, et al. In vivo whole-cortex marker of excitation-inhibition ratio indexes cortical maturation and cognitive ability in youth. *Proceedings of the National Academy of Sciences*, 121(23):e2318641121, 2024.

[39] Pablo Barttfeld, Lynn Uhrig, Jacobo D Sitt, Mariano Sigman, Béchir Jarraya, and Stanislas Dehaene. Signature of consciousness in the dynamics of resting-state brain activity. *Proceedings of the National Academy of Sciences*, 112(3):887–892, 2015.

[40] M. Hashemi, A. Hutt, H. Darren, and J. Sleigh. Anesthetic action on the transmission delay between cortex and thalamus explains the beta-buzz observed under propofol anesthesia. *PLOS ONE*, 12(6):1–29, 06 2017.

[41] Andrea I Luppi, Joana Cabral, Rodrigo Cofre, Pedro AM Mediano, Fernando E Rosas, Abid Y Qureshi, Amy Kuceyeski, Enzo Tagliazucchi, Federico Raimondo, Gustavo Deco, et al. Computational modelling in disorders of consciousness: Closing the gap towards personalised models for restoring consciousness. *NeuroImage*, 275:120162, 2023.

[42] Yonatan Sanz Perl, Carla Pallavicini, Juan Piccinini, Athena Demertzi, Vincent Bonhomme, Charlotte Martial, Rajanikant Panda, Naji Alnagger, Jitka Annen, Olivia Gosseries, et al. Low-dimensional organization of global brain states of reduced consciousness. *Cell Reports*, 42(5), 2023.

[43] Viktor Jirsa, Olaf Sporns, Michael Breakspear, Gustavo Deco, and Anthony Randal McIntosh. Towards the virtual brain: network modeling of the intact and the damaged brain. *Archives italiennes de biologie*, 148(3):189–205, 2010.

[44] Paula Sanz Leon, Stuart Knock, M. Woodman, Lia Domide, Jochen Mersmann, Anthony McIntosh, and Viktor Jirsa. The virtual brain: a simulator of primate brain network dynamics. *Frontiers in Neuroinformatics*, 7:10, 2013.

[45] Viktor Jirsa, Huifang Wang, Paul Triebkorn, Meysam Hashemi, Jayant Jha, Jorge Gonzalez-Martinez, Maxime Guye, Julia Makhalova, and Fabrice Bartolomei. Personalised virtual brain models in epilepsy. *The Lancet Neurology*, 22(5):443–454, 2023.

[46] Huifang E Wang, Paul Triebkorn, Martin Breyton, Borana Dollomaja, Jean-Didier Lemarechal, Spase Petkoski, Pierpaolo Sorrentino, Damien Depannemaecker, Meysam Hashemi, and Viktor K Jirsa. Virtual brain twins: from basic neuroscience to clinical use. *National Science Review*, 11(5):nwae079, 2024.

[47] V.K. Jirsa, T. Proix, D. Perdikis, M.M. Woodman, H. Wang, J. Gonzalez-Martinez, C. Bernard, C. Bénar, M. Guye, P. Chauvel, and F. Bartolomei. The virtual epileptic patient: Individualized whole-brain models of epilepsy spread. *NeuroImage*, 145:377 – 388, 2017. Individual Subject Prediction.

[48] Timothée Proix, Fabrice Bartolomei, Maxime Guye, and Viktor K. Jirsa. Individual brain structure and modelling predict seizure propagation. *Brain*, 140(3):641–654, 2017.

[49] P Sorrentino, A Pathak, A Ziaeemehr, E Troisi Lopez, L Cipriano, A Romano, M Sparaco, M Quarantelli, A Banerjee, G Sorrentino, et al. The virtual multiple sclerosis patient. *iScience*, 27(7), 2024.

[50] C. Mazzara, A. Ziaeemehr, E. Troisi Lopez, L. Cipriano, M. Angiolelli, M. Sparaco, M. Quarantelli, C. Granata, G. Sorrentino, M. Hashemi, V. Jirsa, and P. Sorrentino. Mapping brain lesions to conduction delays: The next step for personalized brain models in multiple sclerosis. *Human Brain Mapping*, 46(7):e70219, 2025.

[51] Bahar Hazal Yalcinkaya, Abolfazl Ziaeemehr, Jan Fousek, Meysam Hashemi, Mario Lavanga, Ana Solodkin, Randy McIntosh, Viktor Jirsa, and Spase Petkoski. Personalized virtual brains of alzheimer's disease link dynamical biomarkers of fmri with increased local excitability. *medRxiv*, pages 2023–01, 2023.

[52] Yonatan Sanz Perl, Sol Fittipaldi, Cecilia Gonzalez Campo, Sebastián Moguilner, Josephine Cruzat, Matias E Fraile-Vazquez, Rubén Herzog, Morten L Kringelbach, Gustavo Deco, Pavel Prado, et al. Model-based whole-brain perturbational landscape of neurodegenerative diseases. *Elife*, 12:e83970, 2023.

[53] Kyesam Jung, Esther Florin, Kaustubh R Patil, Julian Caspers, Christian Rubbert, Simon B Eickhoff, and Oleksandr V Popovych. Whole-brain dynamical modelling for classification of parkinson's disease. *Brain communications*, 5(1):fcac331, 2023.

[54] Marianna Angiolelli, Damien Depannemaecker, Hasnae Agouram, Jean Regis, Romain Carron, Marmaduke Woodman, Letizia Chiodo, Paul Triebkorn, Abolfazl Ziaeemehr, Meysam Hashemi, et al. The virtual parkinsonian patient. *npj Systems Biology and Applications*, 11(1):40, 2025.

[55] Gustavo Deco and Morten L Kringelbach. Great expectations: using whole-brain computational connectomics for understanding neuropsychiatric disorders. *Neuron*, 84(5):892–905, 2014.

[56] Behzad Iravani, Artin Arshamian, Peter Fransson, and Neda Kaboodvand. Whole-brain modelling of resting state fmri differentiates adhd subtypes and facilitates stratified neuro-stimulation therapy. *Neuroimage*, 231:117844, 2021.

[57] Anna Letizia Allegra Mascaro, Egidio Falotico, Spase Petkoski, Maria Pasquini, Lorenzo Vannucci, Núria Tort-Colet, Emilia Conti, Francesco Resta, Cristina Spalletti, Shravan Tata Ramalingasetty, et al. Experimental and computational study on motor control and recovery after stroke: toward a constructive loop

between experimental and virtual embodied neuroscience. *Frontiers in systems neuroscience*, 14:31, 2020.

[58] Sebastian Idesis, Chiara Favaretto, Nicholas V Metcalf, Joseph C Griffis, Gordon L Shulman, Maurizio Corbetta, and Gustavo Deco. Inferring the dynamical effects of stroke lesions through whole-brain modeling. *NeuroImage: Clinical*, 36:103233, 2022.

[59] Giovanni Rabuffo, Houefa-Armelle Lokossou, Zengmin Li, Abolfazl Ziaee-Mehr, Meysam Hashemi, Pascale P. Quilichini, Antoine Ghestem, Ouafae Arab, Monique Esclapez, Parul Verma, Ashish Raj, Alessandro Gozzi, Pierpaolo Sorrentino, Kai-Hsiang Chuang, Teodora-Adriana Perles-Barbacaru, AngÃ¨le Viola, Viktor K. Jirsa, and Christophe Bernard. Mapping global brain reconfigurations following local targeted manipulations. *Proceedings of the National Academy of Sciences*, 122(16):e2405706122, 2025.

[60] Jin Liu, Min Li, Wei Lan, Fang-Xiang Wu, Yi Pan, and Jianxin Wang. Classification of alzheimer's disease using whole brain hierarchical network. *IEEE/ACM transactions on computational biology and bioinformatics*, 15(2):624–632, 2016.

[61] Gustavo Patow, Leon Stefanovski, Petra Ritter, Gustavo Deco, Xenia Kobeleva, and Alzheimer's Disease Neuroimaging Initiative. Whole-brain modeling of the differential influences of amyloid-beta and tau in alzheimer's disease. *Alzheimer's Research & Therapy*, 15(1):210, 2023.

[62] Timothée Proix, Viktor K Jirsa, Fabrice Bartolomei, Maxime Guye, and Wilson Truccolo. Predicting the spatiotemporal diversity of seizure propagation and termination in human focal epilepsy. *Nature Communications*, 9(1):1088, 2018.

[63] Joana Cabral, Francesca Castaldo, Jakub Vohryzek, Vladimir Litvak, Christian Bick, Renaud Lambiotte, Karl Friston, Morten L Kringelbach, and Gustavo Deco. Metastable oscillatory modes emerge from synchronization in the brain spacetime connectome. *Communications Physics*, 5(1):1–13, 2022.

[64] Yuanzhe Liu, Caio Seguin, Sina Mansour, Stuart Oldham, Richard Betzel, Maria A Di Biase, and Andrew Zalesky. Parameter estimation for connectome generative models: Accuracy, reliability, and a fast parameter fitting method. *Neuroimage*, 270:119962, 2023.

[65] Carl-Magnus Svensson, Stephen Coombes, and Jonathan Westley Peirce. Using Evolutionary Algorithms for Fitting High-Dimensional Models to Neuronal Data. *Neuroinformatics*, 10(2):199–218, 2012.

[66] Meysam Hashemi, Axel Hutt, Laure Buhry, and Jamie Sleigh. Optimal model parameter estimation from eeg power spectrum features observed during general anesthesia. *Neuroinformatics*, 16(2):231–251, Apr 2018.

[67] Meysam Hashemi, Anirudh N Vattikonda, Viktor Sip, Sandra Diaz-Pier, Alexander Peyser, Huifang Wang, Maxime Guye, Fabrice Bartolomei, Marmaduke M Woodman, and Viktor K Jirsa. On the influence of prior information evaluated by fully bayesian criteria in a personalized whole-brain model of epilepsy spread. *PLoS computational biology*, 17(7):e1009129, 2021.

[68] Meysam Hashemi, Anirudh N Vattikonda, Jayant Jha, Viktor Sip, Marmaduke M Woodman, Fabrice Bartolomei, and Viktor K Jirsa. Amortized bayesian inference on generative dynamical network models of epilepsy using deep neural density estimators. *Neural Networks*, 163:178–194, 2023.

[69] Meysam Hashemi, Abolfazl Ziaeemehr, Marmaduke M Woodman, Jan Fousek, Spase Petkoski, and Viktor K Jirsa. Simulation-based inference on virtual brain models of disorders. *Machine Learning: Science and Technology*, 5(3):035019, jul 2024.

[70] 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.

[71] Pedro J Gonçalves, 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.

[72] Andrew Gelman, Christian Robert, Nicolas Chopin, and Judith Rousseau. *Bayesian Data Analysis*. CRC Press, 1995.

[73] Samuel Gershman and Noah Goodman. Amortized inference in probabilistic reasoning. In *Proceedings of the annual meeting of the cognitive science society*, volume 36, 2014.

[74] M. Betancourt and M. Girolami. Hamiltonian monte carlo for hierarchical models. *arXiv:1312.0906*, 2013.

[75] M. Betancourt, S. Byrne, S. Livingstone, and M. Girolami. The geometric foundations of hamiltonian monte carlo. *arXiv:1410.5110*, 2014.

[76] Meysam Hashemi, AN Vattikonda, Viktor Sip, Maxime Guye, Fabrice Bartolomei, Michael Marmaduke Woodman, and Viktor K Jirsa. The Bayesian virtual epileptic patient: A probabilistic framework designed to infer the spatial map of epileptogenicity in a personalized large-scale brain model of epilepsy spread. *NeuroImage*, 217:116839, 2020.

[77] Scott A Sisson, Yanan Fan, and Mark M Tanaka. Sequential monte carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765, 2007.

[78] Mark A Beaumont, Jean-Marie Cornuet, Jean-Michel Marin, and Christian P Robert. Adaptive approximate bayesian computation. *Biometrika*, 96(4):983–990, 2009.

[79] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. *Advances in Neural Information Processing Systems*, 2017.

[80] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows. *Advances in Neural Information Processing Systems*, 32:7511–7522, 2019.

[81] Nina Baldy, Martin Breyton, Marmaduke M Woodman, Viktor K Jirsa, and Meysam Hashemi. Inference on the macroscopic dynamics of spiking neurons. *Neural Computation*, pages 1–43, 08 2024.

[82] B.H. Jansen and V.G. Rit. Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biol. Cybern*, 73:357–366, 1995.

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

[84] Anton A Selivanov, Judith Lehnert, Thomas Dahms, Philipp Hövel, Alexander L Fradkov, and Eckehard Schöll. Adaptive synchronization in delay-coupled networks of stuart-landau oscillators. *Physical Review E–Statistical, Nonlinear, and Soft Matter Physics*, 85(1):016201, 2012.

[85] Kong-Fatt Wong and Xiao-Jing Wang. A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience*, 26(4):1314–1328, 2006.

[86] Gustavo Deco, Adrián Ponce-Alvarez, Dante Mantini, Gian Luca Romani, Patric Hagmann, and Maurizio Corbetta. Resting-state functional connectivity emerges from structurally and dynamically shaped slow linear fluctuations. *Journal of Neuroscience*, 33(27):11239–11252, 2013.

[87] Benoit Duchet, Filippo Ghezzi, Gihan Weerasinghe, Gerd Tinkhauser, Andrea A Kühn, Peter Brown, Christian Bick, and Rafal Bogacz. Average beta burst duration profiles provide a signature of dynamical changes between the on and off medication states in parkinson's disease. *PLoS computational biology*, 17(7):e1009116, 2021.

[88] James J Sermon, Maria Olaru, Juan Anso, Stephanie Cernera, Simon Little, Maria Shcherbakova, Rafal Bogacz, Philip A Starr, Timothy Denison, and Benoit Duchet. Sub-harmonic entrainment of cortical gamma oscillations to deep brain stimulation in parkinson's disease: model based predictions and validation in three human subjects. *Brain Stimulation*, 16(5):1412–1424, 2023.

[89] Hugh R Wilson and Jack D Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12(1):1–24, 1972.

[90] 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.

[91] Andreas Daffertshofer and Bernadette CM van Wijk. On the influence of amplitude on the connectivity between phases. *Frontiers in neuroinformatics*, 5:6, 2011.

[92] Ben H Jansen and Vincent G Rit. Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biological cybernetics*, 73(4):357–366, 1995.

[93] Olivier David, Stefan J. Kiebel, Lee M. Harrison, Jérémie Mattout, James M. Kilner, and Karl J. Friston. Dynamic causal modeling of evoked responses in eeg and meg. *NeuroImage*, 30(4):1255 – 1272, 2006.

[94] R.J. Moran, S.J. Kiebel, K.E. Stephan, R.B. Reilly, J. Daunizeau, and K.J. Friston. A neural mass model of spectral responses in electrophysiology. *NeuroImage*, 37(3):706 – 720, 2007.

[95] Fabrice Wendling, Fabrice Bartolomei, Jean-Jcques Bellanger, and Patrick Chauvel. Interpretation of interdependencies in epileptic signals using a macroscopic physiological model of the eeg. *Clinical neurophysiology*, 112(7):1201–1218, 2001.

[96] Sheida Kazemi and Yousef Jamali. On the influence of input triggering on the dynamics of the jansen-rit oscillators network. *arXiv preprint arXiv:2202.06634*, 2022.

[97] Gustavo Deco, Morten L Kringelbach, Viktor K Jirsa, and Petra Ritter. The dynamics of resting fluctuations in the brain: metastability and its dynamical cortical core. *Scientific reports*, 7(1):3095, 2017.

[98] Spase Petkoski and Viktor K Jirsa. Transmission time delays organize the brain network synchronization. *Philosophical Transactions of the Royal Society A*, 377(2153):20180132, 2019.

[99] Jean-Didier Lemaréchal, Maciej Jedynak, Lena Trebaul, Anthony Boyer, François Tadel, Manik Bhattacharjee, Pierre Deman, Viateur Tuyisenge, Leila Ayoubian, Etienne Hugues, et al. A brain atlas of axonal and synaptic delays based on modelling of cortico-cortical evoked potentials. *Brain*, 145(5):1653–1667, 2022.

[100] Pierpaolo Sorrentino, Spase Petkoski, Maddalena Sparaco, E Troisi Lopez, Elisabetta Signoriello, Fabio Baselice, Simona Bonavita, Maria Agnese Pirozzi, Mario Quarantelli, Giuseppe Sorrentino, et al. Whole-brain propagation delays in multiple sclerosis, a combined tractography-magnetoencephalography study. *Journal of Neuroscience*, 42(47):8807–8816, 2022.

[101] Agnes P Funk and Charles M Epstein. Natural rhythm: evidence for occult 40 hz gamma oscillation in resting motor cortex. *Neuroscience letters*, 371(2-3):181–184, 2004.

[102] György Buzsáki and Xiao-Jing Wang. Mechanisms of gamma oscillations. *Annual review of neuroscience*, 35(1):203–225, 2012.

[103] Alfons Schnitzler and Joachim Gross. Normal and pathological oscillatory communication in the brain. *Nature reviews neuroscience*, 6(4):285–296, 2005.

[104] Maria Luisa Saggio, Dakota Crisp, Jared M Scott, Philippa Karoly, Levin Kuhlmann, Mitsuyoshi Nakatani, Tomohiko Murai, Matthias DÃ¼mpelmann, Andreas Schulze-Bonhage, Akio Ikeda, Mark Cook, Stephen V Gliske, Jack Lin, Christophe Bernard, Viktor Jirsa, and William C Stacey. A taxonomy of seizure dynamotypes. *eLife*, 9:e55632, jul 2020.

[105] Herman Haken. Synergetics. *Physics Bulletin*, 28(9):412, 1977.

[106] Viktor K Jirsa and Hermann Haken. A derivation of a macroscopic field theory of the brain from the quasi-microscopic neural dynamics. *Physica D: Nonlinear Phenomena*, 99(4):503–526, 1997.

[107] Timothée Proix, Fabrice Bartolomei, Patrick Chauvel, Christophe Bernard, and Viktor K. Jirsa. Permittivity coupling across brain regions determines seizure recruitment in partial epilepsy. *Journal of Neuroscience*, 34(45):15009–15021, 2014.

[108] Anthony R. McIntosh and Viktor K. Jirsa. The hidden repertoire of brain dynamics and dysfunction. *Network Neuroscience*, 3(4):994–1008, 2019.

[109] Áine Byrne, Reuben D O'Dea, Michael Forrester, James Ross, and Stephen Coombes. Next-generation neural mass and field modeling. *Journal of neurophysiology*, 123(2):726–742, 2020.

[110] Giovanni Rabuffo, Jan Fousek, Christophe Bernard, and Viktor Jirsa. Neuronal cascades shape whole-brain functional dynamics at rest. *ENeuro*, 8(5), 2021.

[111] M Breyton, J Fousek, G Rabuffo, P Sorrentino, L Kusch, M Massimini, S Petkoski, and V Jirsa. Spatiotemporal brain complexity quantifies consciousness outside of perturbation paradigms. *bioRxiv*, pages 2023–04, 2023.

[112] Jan Fousek, Giovanni Rabuffo, Kashyap Gudibanda, Hiba Sheheitli, Spase Petkoski, and Viktor Jirsa. Symmetry breaking organizes the brain's resting state manifold. *Scientific Reports*, 14(1):31970, 2024.

[113] Karl J Friston, Andrea Mechelli, Robert Turner, and Cathy J Price. Nonlinear responses in fmri: the balloon model, volterra kernels, and other hemodynamics. *NeuroImage*, 12(4):466–477, 2000.

[114] Enrique C.A. Hansen, Demian Battaglia, Andreas Spiegler, Gustavo Deco, and Viktor K. Jirsa. Functional connectivity dynamics: Modeling the switching behavior of the resting state. *NeuroImage*, 105:525 – 535, 2015.

[115] Gustavo Deco, Morten L Kringelbach, Aurina Arnatkeviciute, Stuart Oldham, Kristina Sabaroedin, Nigel C Rogasch, Kevin M Aquino, and Alex Fornito. Dynamical consequences of regional heterogeneity in the brain's transcriptional landscape. *Science Advances*, 7(29):eabf4752, 2021.

[116] Anita Monteverdi, Fulvia Palesi, Michael Schirner, Francesca Argentino, Mariateresa Merante, Alberto Redolfi, Francesca Conca, Laura Mazzocchi, Stefano F Cappa, Matteo Cotta Ramusino, et al. Virtual brain simulations reveal network-specific parameters in neurodegenerative dementias. *Frontiers in Aging Neuroscience*, 15:1204134, 2023.

[117] Pedro Costa Klein, Ulrich Ettinger, Michael Schirner, Petra Ritter, Dan Rujescu, Peter Falkai, Nikolaos Koutsouleris, Lana Kambeitz-Ilankovic, and Joseph Kambeitz. Brain network simulations indicate effects of neuregulin-1 genotype on excitation-inhibition balance in cortical dynamics. *Cerebral Cortex*, 31(4):2013–2025, 2021.

[118] Pedro Costa Klein, Ulrich Ettinger, Michael Schirner, Petra Ritter, Dan Rujescu, Peter Falkai, Nikolaos Koutsouleris, Lana Kambeitz-Ilankovic, and Joseph Kambeitz. Brain network simulations indicate effects of neuregulin-1 genotype on excitation-inhibition balance in cortical dynamics. *Cerebral Cortex*, 31(4):2013–2025, 2021.

[119] Matthew F Glasser and David C Van Essen. Mapping human cortical areas in vivo based on myelin content as revealed by t1-and t2-weighted mri. *Journal of neuroscience*, 31(32):11597–11616, 2011.

[120] Klaas Enno Stephan, Nikolaus Weiskopf, Peter M Drysdale, Peter A Robinson, and Karl J Friston. Comparing hemodynamic models with dcm. *Neuroimage*, 38(3):387–401, 2007.

[121] Klaas Enno Stephan, Lars Kasper, Lee M Harrison, Jean Daunizeau, Hanneke EM den Ouden, Michael Breakspear, and Karl J Friston. Nonlinear dynamic causal models for fmri. *Neuroimage*, 42(2):649–662, 2008.

[122] Rens van de Schoot, Sarah Depaoli, Ruth King, Bianca Kramer, Kaspar Märtens, Mahlet G Tadesse, Marina Vannucci, Andrew Gelman, Duco Veen, Joukje Willemsen, et al. Bayesian statistics and modelling. *Nature Reviews Methods Primers*, 1(1):1–26, 2021.

[123] Jayant Jha, Meysam Hashemi, Anirudh Nihalani Vattikonda, Huifang Wang, and Viktor Jirsa. Fully bayesian estimation of virtual brain parameters with self-tuning hamiltonian monte carlo. *Machine Learning: Science and Technology*, 3(3):035016, 2022.

[124] George Papamakarios, David Sterratt, and Iain Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 837–848. PMLR, 2019.

[125] 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.

[126] Johann Brehmer, Gilles Louppe, Juan Pavez, and Kyle Cranmer. Mining gold from implicit models to improve likelihood-free inference. *Proceedings of the National Academy of Sciences*, 117(10):5242–5249, 2020.

[127] George Papamakarios and Iain Murray. Fast ε-free inference of simulation models with bayesian conditional density estimation. In *Advances in Neural Information Processing Systems*, pages 1028–1036, 2016.

[128] 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.

[129] Qiao Liu, Jiaze Xu, Rui Jiang, and Wing Hung Wong. Density estimation using deep generative neural networks. *Proceedings of the National Academy of Sciences*, 118(15):e2101344118, 2021.

[130] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning (ICML)*, pages 1530–1538. PMLR, 2015.

[131] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.

[132] Ivan Kobyzev, Simon Prince, and Marcus Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[133] D.M. Bates and D.G. Watts. Relative curvature measures of nonlinearity. *J R Stat Soc Ser B (Methodological)*, 42(1):1–25, 1980.

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

[135] Michael Deistler, Jakob H Macke, and Pedro J Gonçalves. Disparate energy consumption despite similar network activity. *bioRxiv*, 2021.

[136] Paul G Constantine. *Active subspaces: Emerging ideas for dimension reduction in parameter studies*. SIAM, 2015.

[137] Caglar Cakan, Nikola Jajcay, and Klaus Obermayer. neurolib: A simulation framework for whole-brain neural mass modeling. *Cognitive Computation*, Oct 2021.

[138] Marcel Stimberg, Romain Brette, and Dan FM Goodman. Brian 2, an intuitive and efficient neural simulator. *elife*, 8:e47314, 2019.

[139] Chaoming Wang, Tianqiu Zhang, Xiaoyu Chen, Sichao He, Shangyang Li, and Si Wu. Brainpy, a flexible, integrative, efficient, and extensible framework for general-purpose brain dynamics programming. *elife*, 12:e86365, 2023.

[140] Marília Barandas, Duarte Folgado, Letícia Fernandes, Sara Santos, Mariana Abreu, Patrícia Bota, Hui Liu, Tanja Schultz, and Hugo Gamboa. Tsfel: Time series feature extraction library. *SoftwareX*, 11:100456, 2020.

[141] Oliver M Cliff, Annie G Bryant, Joseph T Lizier, Naotsugu Tsuchiya, and Ben D Fulcher. Unifying pairwise interactions in complex dynamics. *Nature Computational Science*, 3(10):883–893, 2023.

[142] Ben D Fulcher and Nick S Jones. hctsa: A computational framework for automated time-series phenotyping using massive feature extraction. *Cell systems*, 5(5):527–531, 2017.

[143] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[144] Carl H Lubba, Sarab S Sethi, Philip Knaute, Simon R Schultz, Ben D Fulcher, and Nick S Jones. catch22: Canonical time-series characteristics: Selected through highly comparative time-series analysis. *Data Mining and Knowledge Discovery*, 33(6):1821–1852, 2019.

[145] Paul Triebkorn, Leon Stefanovski, Kiret Dhindsa, Margarita-Arimatea Diaz-Cortes, Patrik Bey, Konstantin Bülau, Roopa Pai, Andreas Spiegler, Ana Solodkin, Viktor Jirsa, et al. Brain simulation augments machine-learning–based classification of dementia. *Alzheimer's & Dementia: Translational Research & Clinical Interventions*, 8(1):e12303, 2022.

[146] Leon Stefanovski, Paul Triebkorn, Andreas Spiegler, Margarita-Arimatea Diaz-Cortes, Ana Solodkin, Viktor Jirsa, Anthony Randal McIntosh, Petra Ritter, and Alzheimer's Disease Neuroimaging Initiative. Linking molecular pathways and large-scale computational modeling to assess candidate disease mechanisms and pharmacodynamics in alzheimer's disease. *Frontiers in computational neuroscience*, 13:54, 2019.

[147] Mark A. Beaumont, Wenyang Zhang, and David J. Balding. Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.

[148] Mark A Beaumont. Approximate bayesian computation in evolution and ecology. *Annual review of ecology, evolution, and systematics*, 41:379–406, 2010.

[149] S.A. Sisson, Y. Fan, and M. Beaumont. *Handbook of Approximate Bayesian Computation*. Chapman & Hall/CRC handbooks of modern statistical methods. CRC Press, Taylor & Francis Group, 2018.

[150] Nina Baldy, Nicolas Simon, Viktor Jirsa, and Meysam Hashemi. Hierarchical bayesian pharmacometrics analysis of baclofen for alcohol use disorder. *Machine Learning: Science and Technology*, 2023.

[151] Jan-Matthis Lueckmann, Giacomo Bassetto, Theofanis Karaletsos, and Jakob H Macke. Likelihood-free inference with emulator networks. In *Symposium on Advances in Approximate Bayesian Inference*, pages 32–53. PMLR, 2019.

[152] Conor Durkan, Iain Murray, and George Papamakarios. On contrastive learning for likelihood-free inference. In *International Conference on Machine Learning*, pages 2771–2781. PMLR, 2020.

[153] W.D. Penny, K.E. Stephan, A. Mechelli, and K.J. Friston. Comparing dynamic causal models. *NeuroImage*, 22(3):1157 – 1172, 2004.

[154] W.D. Penny. Comparing dynamic causal models using aic, bic and free energy. *NeuroImage*, 59(1):319 – 330, 2012. Neuroergonomics: The human brain in action and at work.

[155] Nina Baldy, Marmaduke Woodman, Viktor K Jirsa, and Meysam Hashemi. Dynamic causal modelling in probabilistic programming languages. *Journal of the Royal Society Interface*, 22(227):20240880, 2025.

[156] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. *Advances in neural information processing systems*, 27, 2014.

[157] Carl Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016.

[158] Viktor Sip, Meysam Hashemi, Timo Dickscheid, Katrin Amunts, Spase Petkoski, and Viktor Jirsa. Characterization of regional differences in resting-state fmri with a data-driven network model of brain dynamics. *Science Advances*, 9(11):eabq7547, 2023.

[159] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

[160] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE signal processing magazine*, 35(1):53–65, 2018.