



Adiabatic dynamic causal modelling

Amirhossein Jafarian^{a,b,*}, Peter Zeidman^b, Rob. C Wykes^{c,d}, Matthew Walker^c, Karl J. Friston^b

^a Cambridge Centre for Frontotemporal Dementia and Related Disorders, Department of Clinical Neurosciences, University of Cambridge, UK

^b The Wellcome Centre for Human Neuroimaging, UCL Queen Square Institute of Neurology, UK

^c Department of Clinical & Experimental Epilepsy, UCL Queen Square Institute of Neurology, UK

^d Nanomedicine Lab, University of Manchester, UK

ARTICLE INFO

Keywords:

Dynamic causal modelling
Cross spectral density
Phase transition
Adiabatic approximation
Bayesian model selection
Bayesian model reduction

ABSTRACT

This technical note introduces adiabatic dynamic causal modelling, a method for inferring slow changes in biophysical parameters that control fluctuations of fast neuronal states. The application domain we have in mind is inferring slow changes in variables (e.g., extracellular ion concentrations or synaptic efficacy) that underlie phase transitions in brain activity (e.g., paroxysmal seizure activity). The scheme is efficient and yet retains a biophysical interpretation, in virtue of being based on established neural mass models that are equipped with a slow dynamic on the parameters (such as synaptic rate constants or effective connectivity). In brief, we use an adiabatic approximation to summarise fast fluctuations in hidden neuronal states (and their expression in sensors) in terms of their second order statistics; namely, their complex cross spectra. This allows one to specify and compare models of slowly changing parameters (using Bayesian model reduction) that generate a sequence of empirical cross spectra of electrophysiological recordings. Crucially, we use the slow fluctuations in the spectral power of neuronal activity as empirical priors on changes in synaptic parameters. This introduces a circular causality, in which synaptic parameters underwrite fast neuronal activity that, in turn, induces activity-dependent plasticity in synaptic parameters. In this foundational paper, we describe the underlying model, establish its face validity using simulations and provide an illustrative application to a chemoconvulsant animal model of seizure activity.

1. Introduction

This paper introduces a class of dynamic causal model (DCM) that can be used for characterising slow fluctuations in biophysical parameters that might underlie phase transitions in the brain. This method is based on a separation of temporal scales (Jirsa et al., 1994; Papadopolou et al., 2017; Rosch et al., 2018a; Rosch et al., 2018b; Rosch et al., 2018c; Blenkinsop et al., 2012; Jirsa et al., 2014; Steyn-Ross and Steyn-Ross, 2010; Nevado-Holgado et al., 2012) where fast neuronal fluctuations are generated by slow fluctuations in synaptic parameters and other neurophysiological parameters (e.g., extracellular potassium). DCM then allows one to specify different hypotheses about causal relations between slow biological mechanisms (Papadopolou et al., 2015) and select the most likely model that explains phase transitions in electrophysiological data. The innovation of the DCM introduced here is that the separation of temporal scales is used to introduce a circular causality in which synaptic parameters shape fast neuronal activity, while frequency specific neuronal activity induces plasticity or changes in synaptic parameters. In other words, fast, (spectral) neuronal dynamics are modulated on a slow timescale by drifts

in synaptic parameters, while the spectral characteristics of fast neuronal activity causes slow changes in the parameters. In this paper, we illustrate how this circular causality and implicit separation of temporal scales leads to the spontaneous onset of phase transitions in brain activity and crucially, how this formulation of (patho)physiology can be used as the basis of a relatively straightforward hierarchical DCM, which we refer to as an adiabatic (or A)-DCM. In brief, the mapping from synaptic parameters to fast – within epoch – neuronal activity uses a conventional DCM for cross spectral density (CSD). In A-DCM, slow – between epoch – changes in spectral density are then used as an empirical prior on synaptic parameters to model activity-dependent plasticity.

The motivation for developing A-DCM was to provide people with a relatively straightforward procedure that enables them to evaluate hypotheses about the underlying causes of phase transitions in neuronal activity, in terms of model evidence. For instance, the circular causality between slow parameters and spectral responses of neuronal oscillations could be useful for understanding the relationship between depth of anaesthesia (induced with gradual drug injections such as Propofol to modulate the frequency contents of electrophysiological

* Corresponding author at: Cambridge Centre for Frontotemporal Dementia and Related Disorders, Department of Clinical Neurosciences, University of Cambridge, Herchel Smith Building, Forvie Site, Robinson Way, Cambridge Biomedical Campus, Cambridge CB2 0SZ, UK.

E-mail address: aj631@cam.ac.uk (A. Jafarian).

<https://doi.org/10.1016/j.neuroimage.2021.118243>.

Received 26 August 2020; Received in revised form 3 June 2021; Accepted 4 June 2021

Available online 8 June 2021.

1053-8119/© 2021 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

recordings) and slow dynamics of synaptic efficacy of neuronal populations (Purdon et al., 2013; Hashemi and Hutt, 2016). Adiabatic DCM could be useful for investigating how the slow evolution of ion/synaptic mechanisms engenders brain state transitions, e.g., awake to asleep (Muheim et al., 2019) or pathological states, e.g. seizures (Grenier et al., 2003). Exemplar questions we envisage being addressed using A-DCM include: (i) what sorts of slow biological mechanisms might account for the decline of alpha band power as a marker of progression in Alzheimer's disease? (Li et al., 2020); (ii) what are the links between induced paroxysmal epileptic transitions (3 to 8 Hz) and pathological evolution of potassium level or synaptic transmission mechanisms? (Moody Jr et al., 1974; Du et al., 2016); (iii) what is the role of changes in synaptic efficacy in inhibitory populations and the emergence of high beta activity, as hallmarks of pathophysiology in Parkinson's disease (McCarthy et al., 2011)? We do not address these long-standing questions here. Instead, through a worked example of A-DCM we show how one can (i) establish causal links between biological parameters and spectral responses in generative models of neuroimaging data; (ii) evaluate model evidence and infer parameters from empirical data and; (iii) compare different hypotheses (through Bayesian model reduction and comparison) about how the data were generated.

With regards to extant procedures, A-DCM complements previous approaches to fast-slow modelling of phase transitions (Steyn-Ross and Steyn-Ross, 2010; Coombes and Bressloff, 2005; Baier et al., 2012; Wendling et al., 2016; Schiff, 2012). Conventionally, multiscale models feature fast states that constitute neuronal dynamics, which are (predominantly) regulated by the dynamics of some slow states, such as ion concentrations and synaptic efficacy. The prominent application of multiscale models is to understand the taxonomy and phenomenology of phase transitions in the neuroimaging data, in particular electrophysiological recordings, as reviewed in (Friston, 2014). Some recent (selective) examples of these models are: (i) a neuronal-glia interaction model that emulates epileptic seizures and cortical spreading depression through slow dynamics of potassium and sodium (Wei et al., 2014); (ii) a neural mass model (NMM) augmented with slow dynamics for synaptic efficacy—e.g. as a model of synaptic plasticity (Fung and Robinson, 2014)—to replicate burst suppression in anaesthesia (Liley and Walsh, 2013; Liu and Ching, 2017); (iii) a slow-fast mesoscale model of epileptic seizures that captures the slow dynamics of firing thresholds—to replicate spike rate adaptation—with a NMM (Jafarian et al., 2019b); and (iv) a phenomenological model of epileptic seizures, known as an Epileptor, where the aetiology of epileptic seizures is explained via the evolution of a slow state (Jirsa et al., 2014; El Houssaini et al., 2020). The novelty of adiabatic DCM, as a forward model, is the ability to link slow biological mechanisms with the spectral contents of fast neuronal dynamics. In short, A-DCM was designed to characterise (paroxysmal) transitions formally, by coupling modulatory slow dynamics of ion current/synaptic efficacy and induced spectral responses in mesoscale dynamics.

In terms of model estimation, due to the complexities of multiscale models, inferring the parameters of these models—and comparing models given real electrophysiological recordings—typically incurs a high computational burden (Hashemi et al., 2020; Jafarian et al., 2019b; Jafarian et al., 2019a). In this paper, we show that A-DCM provides a computationally efficient method for parameter estimation—and more importantly calculating model evidence—that can be used to investigate the aetiology of phase transitions from electrophysiological recordings. In particular, model inversion in A-DCM is motivated by Synergetic theory and the Adiabatic approximation (Basar et al., 2012; Haken, 1977; Jirsa and Haken, 1997). This formulation assumes that for each value of slowly varying parameters, the fast-neuronal states attain nonequilibrium steady-state that is expressed in terms of spectral contents. This assumption leads to a relatively straightforward approach to inferring model parameters from empirical data. Furthermore, A-DCM enables one to compare different models or hypotheses about underlying generators of paroxysmal transitions.

In summary, we have introduced two novel contributions in this paper. First, in contrast to the existing slow-fast formulation of phase transitions in the brain (e.g. Epileptor (Jirsa et al., 2014) or slow-fast neural mass model (Jafarian et al., 2019b)), we introduced the modulation of parameters with respect to the frequency contents of fast membrane potentials (both in terms of forward simulations of the model and the estimation procedure). We therefore provided a platform that could be useful for clinical research (in particular, regarding the circular relationship between spectral responses in electrophysiological recordings and their associated underlying biological causes. Second, by employing the spectral contents of the data as a prior when inverting/estimating a neuronal model, we provided an explicit model of circular causality. In addition, we established a procedure for data driven model construction between slow states' dynamics and fast neuronal responses, which has not been considered previously (e.g., we employed a polynomial linear model to relate slow states and the evolution of the spectral contents of the data).

This paper comprises four sections. In the next section, we review the theoretical tenets of A-DCM. In section three, we first provide an illustrative example of the basic ideas in terms of a forward simulation of the generative model. This example illustrates how changes in model parameters can induce transitions in brain activity and accompanying spectral responses. We then use these forward simulations to create a hierarchical generative model (Friston et al., 2016) that can be inferred from empirical data. The third section presents some applications of A-DCM—using Bayesian model reduction—to characterise the underlying causes of seizures. In this section, we generated synthetic data that undergoes transition into and out of seizures. Then, by performing a second level analysis, under a hierarchical model of slow changes in DCM parameters (Friston et al., 2015), we show how one can answer some fundamental questions concerning the genesis of epileptic seizures. This section can be read as establishing the face validity of the procedure. In the fourth section, we apply model inversion to an empirical electrophysiological recordings from an animal model of epilepsy to provide an illustrative (worked) example of this type of analysis. We conclude with a discussion of the limitations and potential applications of adiabatic DCM.

2. Theory

This section provides a brief review of dynamic causal modelling. Then, we build on this to introduce the A-DCM methodology that uses a hierarchical model, in which slow changes in DCM parameters at the second level are coupled to spectral responses at the first.

2.1. Dynamic causal modelling

Dynamical causal modelling is the estimation of biologically informed models of neuroimaging data using variational Bayesian methods (Friston et al., 2007; Friston et al., 2003; Friston et al., 2008; Friston et al., 2019). DCM was pioneered nearly 20 years ago and since been used to infer the biological mechanisms generating neuroimaging data (Jafarian et al., 2019c; Friston, 2011; Friston et al., 2011; Penny et al., 2011; van Wijk et al., 2018). In DCM, a posterior probability density over model parameters, as well as the evidence for a model (for any given empirical data), are inferred through optimisation of an objective function called the variational free energy. This objective is known as an evidence lower bound (ELBO) in machine learning and provides a computationally efficient approximation to the (log) model evidence or marginal likelihood. This optimisation is performed under the Laplace approximation (i.e., probability densities are approximated using Gaussian distributions) (Friston et al., 2007; Friston et al., 2008; Beal, 2003; Zeidman et al., 2019a) using a gradient ascent on variational free energy. This is known as Variational Laplace. In DCM, the model evidence associated with different hypotheses (i.e., models) of the same data are compared using Bayesian model selection and comparison to identify

Table 1
Parameters of the neuronal model (see also Fig. 2).

	Description	Parameterisation	Prior
T	Postsynaptic rate constant	$\exp(\theta_c) \cdot T_i T = [256, 128, 16, 32]$	$p(\theta_c) = N(0, 1/16)$
g	Intrinsic connectivity between populations i and k in each region	$\exp(\theta_a) \cdot g$	$p(\theta_a) = N(0, 1/16)$
G	Self-inhibitory connection	$\exp(\theta_a) \cdot G$	$p(\theta_a) = N(0, 1/16)$

Table 2
Glossary of variables and expressions.

Variable	Description
u	Exogenous input
x	The i -th (neuronal) state in region j ; e.g., mean depolarisation of a neuronal population
$\sigma(x)$	The neuronal firing rate – a sigmoid squashing function of depolarisation
L	Lead field vector mapping from (neuronal) states to measured (electrophysiological) responses
$g_u(\omega), g_o(\omega), g_y(\omega)$	Spectral density of (neuronal) state fluctuations, observation error and ensuing measurement, respectively
$\partial_x f$	System Jacobian or derivative of system flow with respect to (neuronal) states
$k(t) = FT[K(\omega)]$	First order kernel mapping from inputs to responses; c.f., an impulse response function of time. This is the Fourier transform of the transfer function
$K(\omega) = FT[k(t)]$	Transfer function of frequency, modulating the power of endogenous neuronal fluctuations to produce a (cross spectral density) response. This is the Fourier transform of the kernel

the best explanation for the data at hand (Kass and Raftery, 1995). A recently developed Bayesian model reduction technique, which we will leverage here, opens a new avenue for rapidly comparing the evidence for models specified in terms of their priors (Friston et al., 2018; Friston and Penny, 2011; Friston et al., 2016; Zeidman et al., 2019b). A greedy (i.e., depth-first) search is performed in the BMR method that iteratively switches off mixtures of parameters (Friston et al., 2018). Each mixture is removed when switching it off causes no reduction in free energy (i.e., when the parameters were only contributing to complexity, without increasing accuracy). This search is highly efficient and takes seconds on a standard desktop computer. For more details, please see (Friston et al., 2018; Jafarian et al., 2019c).

2.2. Adiabatic dynamic causal model (A-DCM)

2.2.1. Theoretical foundation

The starting point for A-DCM is a separation of temporal scales into fast (neuronal) and slow fluctuations (Basar et al., 2012). Our main assumption here is that the brain can be treated as an excitable medium, where neuronal dynamics operate around a sequence of fixed points, where oscillatory dynamics result from the (neuronal) filtering of endogenous random fluctuations (Moran et al., 2007; Moran et al., 2013; Friston et al., 2012). The fixed points then change slowly to generate changes in cross spectral density over time (please also see Appendix C for further discussion) (Table 1 and 2).

The slow and fast scale separation, which is known as an adiabatic approximation (Haken, 1977), leads naturally to a mean field approximation when building generative models for data analysis. Under a mean field approximation, the posterior over unknown variables is approximated with the product of (marginal) posteriors (Friston et al., 2007). In this instance, the adiabatic approximation enables us to estimate slowly changing synaptic (i.e. connectivity) parameters and hyperparameters, under the assumption that they are conditionally independent. The mean field approximation requires us to write down the equations of motion for the parameters or connection strengths. Effectively, this entails specifying a model of synaptic plasticity. We therefore consider a paroxysmal (e.g., epileptogenic) perturbation or phase transition that is mediated by neuronal plasticity (e.g., spike rate adaptation or associative plasticity) to reset the fast-neuronal dynamics – and their expression in spectral responses (please also see Appendix C for further discussion).

In A-DCM, neuronal dynamics are generated in terms of their second order (statistical) moments, via complex cross spectra (Friston et al.,

2012). These then generate observed cross spectra in the sensor domain (via a standard electromagnetic forward model). In detail, let us write the temporal dynamics of neuronal states, x , driven by random fluctuations, u , for a given set of parameters θ as follows:

$$\dot{x} = f(x, \theta) + u \quad (1)$$

In Eq. 1, the cross spectral content of the random fluctuation, $g_u(\omega, \theta) = FT(E[u(t), u(t - \tau)])$, is modelled as (structured) pink noise (Friston et al., 2012; Moran et al., 2007; Moran et al., 2013). The assumption underlying DCM for CSD is that random fluctuations induce oscillations around some fixed point (Da Silva et al., 1974; Friston et al., 2012). In other words, using DCM for CSD, we treat electrophysiological recording, y , as a neuronally filtered version of endogenous noise (Da Silva et al., 1974). Under a fixed point assumption, the neuronal dynamics can be well approximated by the first order linearised neuronal model $\dot{x} = (\nabla_{x^*} f)x + u$ (where ∇_{x^*} denotes the Jacobian of neuronal states at x^*) (David et al., 2006). The linearised equation of neuronal states is used to obtain a semi analytic expression of the transfer function (i.e., first order Volterra kernel) with an impulse response $k(\tau, \theta) = \exp(\tau \cdot \nabla_x f(x, \theta))$. The spectral response of the transfer function can in turn be expressed as follows:

$$K(\omega, \theta) = FT(\exp \tau \cdot \nabla_x f(x, \theta)) \quad (2)$$

The neuronal source response, $g_x(\omega)$, of the noise driven model is given by:

$$g_x(\omega) = K(\omega, \theta) \cdot g_u(\omega, \theta) \cdot K(\omega, \theta)^T + g_o(\omega, \theta) \quad (3)$$

In Eq. 3, $g_o(\omega, \theta)$ represents random fluctuations due to observation noise (associated with individual channels and common to all channels) (Moran et al., 2007). The spectral response in sensor space is calculated through the forward electromagnetic model, denoted by L . $M(\omega)$, as follows:

$$g_y(\omega) = L \cdot M(\omega) \cdot g_x(\omega) \cdot M^T(\omega) \cdot L^T + \epsilon \quad (4)$$

In Eq. 4, $g_y(\omega)$ are the cross spectral data and ϵ is a random effect due to computing the cross spectra from finite timeseries, which can be inferred using the variational Bayes approach in DCM (please see Appendix A for explanation of this final term in Eq. 4).

In A-DCM, slow dynamics of parameters, θ_t , have the following general form, which models activity-dependent plasticity at time t , via a dependency on neuronal activity:

$$\theta_t = \zeta_\rho(g(\omega, t)) + \epsilon_t \quad (5)$$

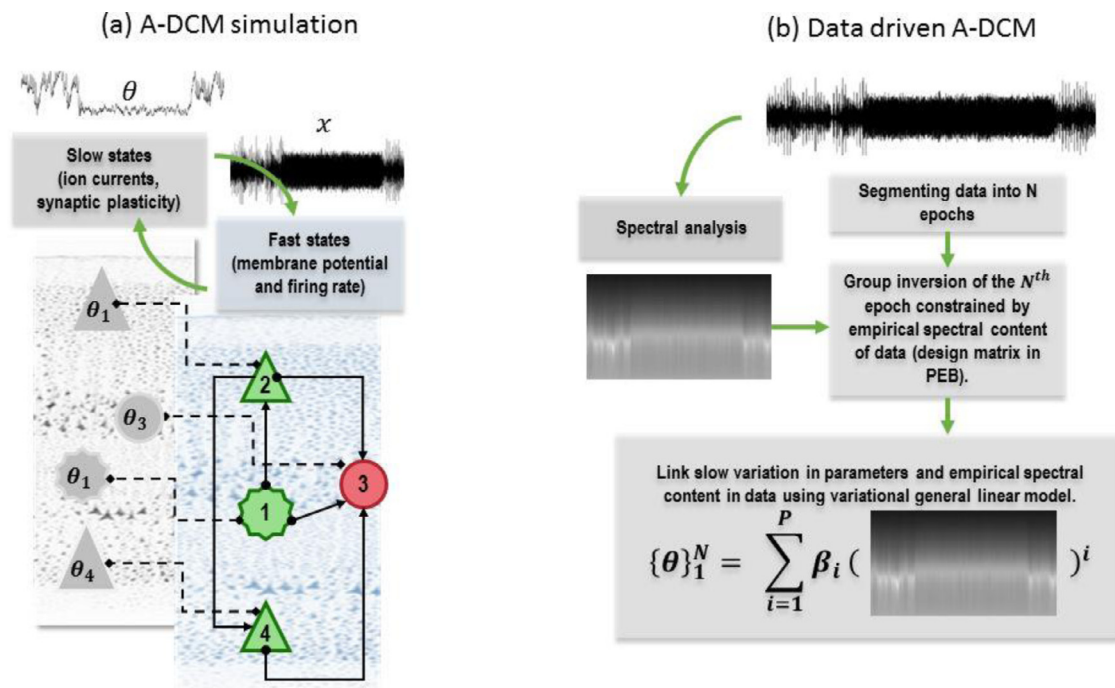


Fig. 1. This figure illustrates forward simulation and estimation procedures in adiabatic DCM. a) Forward simulation procedure for a neuronal model that is equipped with slow states. Here, the value of parameters θ at each point in time induces neuronal responses x . The ensuing induced spectral neuronal responses modulate the value of parameters and so on. In effect, a form of circular causality inherently emerges in A-DCM that accounts for the mechanistic coupling between slow variations of physiological parameters and fast neuronal responses. The four neuronal populations that constitute the canonical microcircuit are 1=spiny stellate cells, 2=superficial pyramidal cells, 3=inhibitory interneurons and 4=deep pyramidal cells (see Fig. 2 for details) and $\theta_1, \dots, \theta_4$ (e.g., self-inhibition) represent parameters of each neuronal populations. Hierarchical inversion in A-DCM rests upon a separation of temporal scales. In this approach, electrophysiological data is divided into segments. Then, hierarchical Bayesian model inversion—using DCM for CSD and PEB—is used to infer time-dependent parameters that are constrained at the second level by the spectral properties of the data. The ensuing estimated parameters are then modelled using a linear in parameters ($\beta_i, i = 1, \dots, P$) polynomial (of order P) GLM, with a design matrix that summarises the spectral features of the data to ensure a coupling between time-dependent parameters and spectral responses in the data. Note that one could use a similar approach to characterise the relation between the slow dynamics of parameters and spectral response of neuronal populations. However, the main aim of A-DCM is to provide a model between physiological parameters and the measurable data at hand.

The term $g(\omega, t)$ in Eq. 5 denotes the spectral activity of neuronal dynamics at epoch t . Here, ζ is an operator (parametrised by unknown parameters ρ) that generates the parameter at time t from the spectral activity of neuronal dynamics, and ϵ_t is additive random effect. One example of the operator S could be the integral of power spectral density (PSD) (Using Parseval's theorem, the variance (average power) of a process can be calculated by integrating the power spectrum over all frequencies: $\text{var}(x) = \int \text{PSD}(x) df$ (Von Storch and Zwiers, 2001)), which engenders a simple form of synaptic plasticity (please see Fig. 1a). Another example of operator S can be concatenated power spectral density of electrophysiological data, e.g., local field potential (LFP) over different epochs. Such definitions for the S operator can be useful when one wants to test for a causal relation between synaptic parameters and frequency specific neuronal activity (i.e., investigating circular causality). In what follows, we will use a general linear model (GLM) based upon regressors (in a design matrix) that encode the expression of particular frequencies. This means that the parameters ρ correspond to the parameters of a GLM (e.g., regression coefficients). Note that the adiabatic approximation allows us to express the slow dynamics as a mapping from the time-dependent spectral density to the parameters (please see Fig. 1b).

In A-DCM, the neural mass model is only used to evaluate the spectral response to some endogenous neuronal fluctuations $u(t)$ that are themselves, parameterised. This enables us to specify a generative model of slow fluctuations in observable (complex cross) spectral density, purely in terms of parametric dynamics, while – at the same time – absorbing a chosen neural mass model into the resulting adiabatic DCM.

2.2.2. Practical implementation

For people familiar with DCM, we employ a standard DCM for cross spectral density (Litvak et al., 2011; Friston et al., 2012) and turn it into a hierarchical state space model by adding dynamics to the parameters. The inversion of this model allows us to estimate the connectivity that best explains empirical cross spectra and the trajectories or dynamical architecture controlling the expression of different synaptic parameters.

Practically, we first divide data into epochs (that may overlap) (please see Fig. 1b). We then use the separation of temporal scales to model slow changes in the parameters from epoch to epoch within a hierarchical or parametric empirical Bayes model (Friston et al., 2015; Friston et al., 2016). In detail, the posteriors over parameters from each epoch are passed to the second level PEB analysis, and slow drifts in the parameters are captured using a general linear model (GLM) (see Papadopoulou et al., 2015). Crucially, the explanatory variables in the GLM comprise the spectral density of neuronal activity in each epoch. This underwrites the circular causality between frequency specific changes in fast, parameter-dependent neuronal responses and the slow, activity-dependent neuronal plasticity. In the examples below, we use the empirical spectra as explanatory variables in the GLM – after some suitable dimension reduction and transformation (Chen et al., 2008; Gavish and Donoho, 2014). This means the empirical spectra are used twice. First, the cross spectrum of each epoch is used as the data feature to fit the parameters of an epoch-specific DCM. Second, between-epoch changes in spectral activity are used to provide empirical priors over changes in these parameters. This dual use of the spectral data is licensed by the separation of temporal scales upon which this DCM rests.

In other words, A-DCM leverages information in the spectral content of each epoch and in spectral changes over epochs.

The reason that this (adiabatic) DCM is efficient is that using a generative model of cross spectra converts a neural state-space model into an instantaneous mapping between the parameters of the model and the expected second order responses over time; namely, the complex cross spectra. The only assumption behind this adiabatic approximation is that the spectral summary of dynamics is sufficient to inform slow changes in parameters. This is a key part of the adiabatic DCM described here: conditioning the parameters of neuronal density dynamics (here, spectral density) means that we can model activity-dependent changes in connectivity and other slowly varying factors that, in turn, shape fast neuronal responses. This means that one can specify an adiabatic DCM to model slow dynamics such as spike rate adaptation, short-term plasticity or, indeed, the target of this work; trajectories in parameter space that engender paroxysmal transitions in neuronal dynamics, e.g., epilepsy. The basic DCM for CSD using this work has been described in many previous applications e.g., (Papadopolou et al., 2017; Rosch et al., 2018a; Rosch et al., 2018b). The key thing that we bring to the table is equipping the model with a second level that is constrained by empirical data features at a slower timescale. In addition, we establish polynomial GLM to model the relationship between slow parameters and spectral contents in data.

To establish the face validity of this kind of model, we will refer to specific empirical data in which seizures were induced chemically. These data and the ensuing seizure activity and now summarised briefly.

2.3. Chemoconvulsant animal model of seizures

Animal experiments were conducted in accordance with the United Kingdom Animal (Scientific Procedures) Act 1986, and approved by the Home Office (license PPL70-13691). Sprague-Dawley rats (8–12 weeks old, 280–330 g; Charles River, UK) were used in this study. All animal experiments were conducted in accordance with the United Kingdom Animal (Scientific Procedures) Act 1986, and approved by the local ethics committee (University College London). Animals were housed on 12 h/12 h dark/light cycle, and food and water were given ad libitum. Animals were group housed and allowed to acclimatise to the new environment for at least 1 week before surgery, and were housed individually after surgery. Rats were anaesthetised using isoflurane (2%) and head-fixed in a stereotaxic frame (Kopf, USA). A small hole was drilled through the skull above the right primary visual cortex (coordinates: 3 mm lateral and 7 mm posterior of bregma and a cannula inserted (Plastics1, USA). During the same surgery, an ECoG transmitter [A3028E-AA, Open Source Instruments] was implanted subcutaneously with a recording electrode wire positioned in the visual cortex. A reference electrode was placed in the contralateral frontoparietal cortex. Animals were single housed in Faraday cages for telemetric ECoG recordings. 7–10 days post-surgery, rats were briefly anaesthetised and 300–400nl of 10mM Picrotoxin administered to layer 5 visual cortex, via the pre-implanted cannula. Immediately after injection rats were removed from the stereotaxic frame and replaced in telemetry. Within a few minutes post-injection small amplitude spikes appeared in the ECoG traces, these evolved over the next 5–10 minutes into large amplitude regular (~1Hz) spikes or poly-spikes. A cyclical pattern of seizures (~53s in duration) and inter-ictal activity was observed for about 2–3 hours before fading away and a resumption of normal brain activity.

3. Validation analyses using simulated data

This paper demonstrates the basic phenomenology that A-DCM is capable of explaining; namely, phase transitions in electrophysiological data. In this section, we will use synthetic data. In subsequent sections, we apply the same methodology to empirical data taken from the mouse model of seizures above, to illustrate the sort of analyses one can perform.

To illustrate A-DCM, we use the canonical microcircuit (CMC) (Bastos et al., 2012; Friston et al., 2019) as a model of electrical activity of a typical cortical column (see Fig. 2 for details). This model has been found to be useful for a range of applications in computational neuroscience, including: predictive coding (Bastos et al., 2012), modelling evoked brain responses (Aukstulewicz and Friston, 2015; Jafarian et al., 2020; Friston et al., 2019) and cross spectral responses (Rosch et al., 2018b; Rosch et al., 2018c) to name a few. The CMC can replicate fast gamma activity of the superficial layers as well as the slow (beta) activity in deep layers in the cortex (Bastos et al., 2012). The CMC comprises four neuronal populations, namely superficial pyramidal cells (layer I of the cortex), excitatory populations (spiny cells in layer IV of the cortex), deep pyramidal cells (layer V of the cortex) and interneuron inhibitory cells. Each population in the CMC receives a firing rate from inter regional and distal neuronal populations, weighted by intrinsic and extrinsic effective connectivity, respectively. Each population converts summed input firing rates to a postsynaptic synaptic response (through convolution of the firing rate and a synaptic impulse response model). The ensuing postsynaptic response generated by each neuronal population is then transformed to mean firing rates (through sigmoid transformation), which is then communicated to other populations via intrinsic and extrinsic efferents. Electrophysiological recordings are generated by the CMC as (mainly) the activity of superficial pyramidal cells plus weighted sum (inferred from data) of the activity of excitatory and deep pyramidal populations.

3.1. Face validity: simulation of beta burst synchronization

3.1.1. Part A: forward simulation of beta bursts

In this section, we validate the basic idea of A-DCM – that a reciprocal coupling between fast neuronal activity and slow drifts in synaptic parameters (e.g., extracellular concentrations or synaptic efficacy) induces phase transitions. This validation rests upon integrating (or solving) coupled differential equations at fast and slow timescales to illustrate that when parameters pass from one regime of parameter space to another, there is a qualitative change in the spectral activity at the fast timescale (Papadopolou et al., 2015).

In this simulation, we illustrate the effect of modulating the self-inhibition of deep and inhibitory cells using the integral of the power spectral density (variance of the signal in the time domain) of their postsynaptic potentials. This can be regarded as a simple model of activity-dependent synaptic plasticity (Fung and Robinson, 2014; Fritschy, 2008). This model was motivated by the fact that there is a relationship between the energy content of neuronal activity and ion dynamics. For example, the activity level of neuronal activity declines after seizures (known as post ictal depression), which is related to the pathological evolution of ion currents (Panayiotopoulos, 2010). Another motivation includes the relationship between energy metabolism due to neuronal activity (caused by ion dynamics) and their level of demand for energy as observed in haemodynamic responses (Rosa et al., 2011; Carmichael et al., 2017).

The ensuing simulated data and evolution of parameters are shown in Fig. 3. In this simulation, the dynamics of parameters induce a high band synchronised beta burst. The dynamics of phase transition in this model are akin to a bifurcation or phase-transition. This is because crossing a phase boundary or ‘separatrix’ in parameter space induces the transitions in mesoscopic activity. Although there is a sudden change in the spectral activity induced by this boundary crossing, the drift of the parameters *per se* is quite smooth and slow (in comparison with the fast neuronal states) (Papadopolou et al., 2017; Rosch et al., 2018a; Rosch et al., 2018c). The synchronised beta burst is a hallmark of movement disorders (McCarthy et al., 2011; Spitzer and Haegens, 2017; Sherman et al., 2016), which may also be induced by drugs/interventions (Rodriguez et al., 2004; Shin et al., 2017) or during memory retrieval (Jansen et al., 2011). Note that the simulated seizure activity in Fig. 3 is entirely self-organised. In other words, the drifts

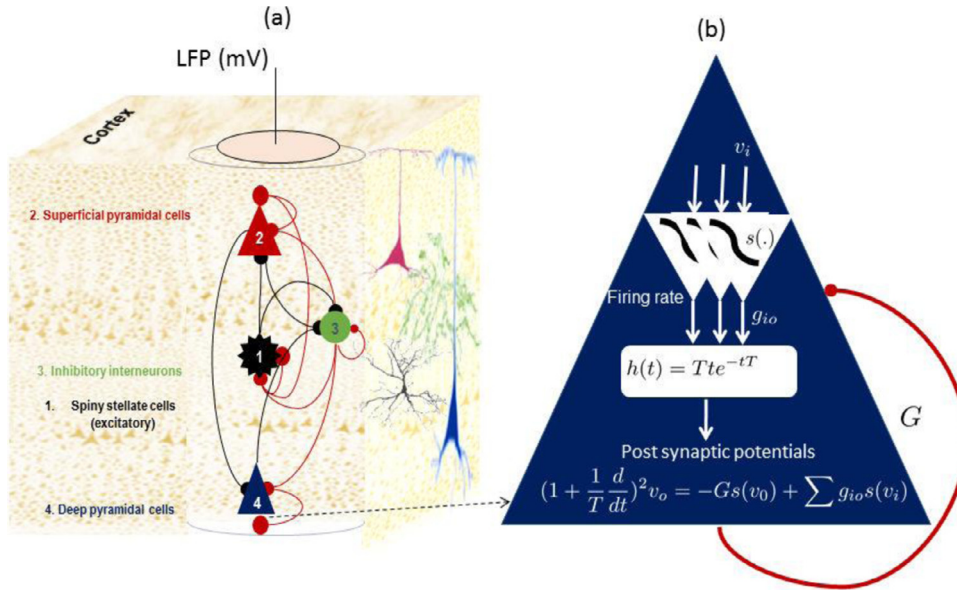


Fig. 2. Canonical microcircuit (CMC). a) A patch of cortex divided into cortical columns. Electrical activity of each cortical column can be captured by electrophysiological recording, e.g., ECoG. Each cortical column is divided into several layers (here 3), each of which is modelled by one population of neurons. Superficial and deep pyramidal cells are in the superficial and deep layers, respectively. Excitatory interneurons (spiny stellate cells) are located in layer four – labelled 1 in the figure. Inhibitory interneurons are distributed across all layers and are modelled using one population that interacts with all other populations. b) The mean electrical activity of each neuronal population is derived using mean field theory using two conversion operators. The postsynaptic potentials, x_i are transformed through a sigmoid nonlinearity, $s(\cdot)$, to generate a firing rate (weighted by connectivity constant g). The ensuing firing rate is then converted to postsynaptic potentials, x_o through the linear response of synapses (parametrised by rate constant T). In addition, each population is equipped with a

self-inhibition connection G (illustrated as short curved red lines) which assures neuronal homeostasis, i.e., in the absence of neuronal input, the activity of neurons rest at an equilibrium. Please see [Appendix D](#) for detailed formulation of this model.

in self-inhibition were driven by activity-dependent plasticity; here, the overall power in neuronal activity of the respective neuronal populations.

The example in [Fig. 3](#) was based upon solving differential equations for fast and slow dynamics in time. To examine the equivalent characterisation in frequency space, we evaluated the transfer function of the CMC model, given the trajectory of parameters in [Fig. 4](#). In detail, we took the trajectory of parameters from [Fig. 3-c](#), and calculated the induced spectral changes through simulation of the transfer function for each sample point in the data. This first-order approximation to the non-linear solution in [Fig. 4](#) confirms that phase transitions simulated above can be attributed to the evolution of parameters. Furthermore, it licenses the use of spectral data features for inversion of a DCM for cross spectral density (CSD) data. We pursue model inversion in the next section.

3.1.2. Part B: inferring adiabatic dynamics

The previous section established the face validity of the generative model in terms of being able to generate plausible phase transitions. In this section, we turn to the face validation of inference or model inversion. In brief, we now try to recover the slow changes in synaptic parameters given the (synthetic) electrophysiological data in the previous section. To infer dynamics of the parameters, first we divide the simulated data above into overlapped epochs (sliding windows that span all samples in the data with the duration two seconds). Then, we implement A-DCM by using the standard (canonical microcircuit) DCM for CSD to explain the cross spectral density of each epoch, furnishing a set of parameters for each window ([Moran et al., 2007](#); [Bastos et al., 2015](#)). The ensuing predictions of spectral responses and estimated self-inhibition of deep pyramidal cell are shown in [Fig. 5](#).

Next, to leverage the circular causality between slow and fast timescales, we inverted a general linear model (GLM) of the parameters from each epoch. The GLM we use for this example is a polynomial expansion of the spectral power:

$$\theta = \beta_1 (g(y)^1) + \beta_2 (g(y)^2) + \dots + \beta_k (g(y)^p) + \varepsilon \quad (6)$$

In [Eq. 6](#), β_i s are unknown coefficients in the GLM, $g(y)^i$ is the restricted power spectral density (to some frequency bins) of the observation data, where i is an exponent that can range from one to an unknown value p . The frequency bins of interest, here, are dominant spectral peaks (which can be identified automatically via a singular value decomposition) before, during and after phase transitions in the observed signal.

To define an optimal value for p , we gradually increased it from one, and compared the associated free energy (i.e., log evidence). At some point during the increase of the polynomial order, the free energy (which is the accuracy minus the complexity of the model) starts to decrease due to overfitting ([Bishop, 2006](#)). The implicit Bayesian model comparison enables one to find the polynomial order p that has the greatest evidence, i.e., the best balance between accuracy and complexity. The results of this model comparison are shown in [Fig. 6](#).

Note that by construction, the spectral data features are both caused by the parameters (at the first or fast level) and also cause the parameters (at the slow or second level). The advantage of this unusual but straightforward construction is that one can test the hypothesis that particular frequencies are responsible for increasing or decreasing particular parameters on a slow timescale. It is this influence of frequency-specific modulation on synaptic parameters that stands in for the reciprocal coupling between fast and slow timescales evinced in the first section.

3.2. Face validity: model comparison

In this section, we present a face validation of the Bayesian model comparison and reduction procedures in A-DCM that can be used to test different hypotheses about how slow (synaptic) changes give rise to spontaneous paroxysmal transitions. To generate synthetic data, we specified trajectories of the self-inhibition of inhibitory and deep pyramidal cells, to produce a characteristic change in spectral density in electrophysiological data (that stands in for the onset of seizure activity). The timeseries solution of the stochastic difference/differential equations generating parameters and simulated data are shown in [Fig. 7](#). In this simulation, pathological transition is due to changes in self-inhibition of the inhibitory and deep pyramidal populations. Interestingly, we observe hysteresis phenomena in the behaviour of the model. This is supported by the fact that when self-inhibition of deep pyramidal and inhibitory cells are increased/decreased, respectively, the model undergoes high frequency activity, which is known to be one of the hallmarks of onsets of spontaneous seizures ([Traub et al., 2001](#); [Grasse et al., 2013](#)). Importantly, when the parameters change in the other direction, the model generates epileptic spikes.

Next, to infer the parameters over 21 predefined epochs, we used DCM for CSD. The reader should note that due to the nonlinear nature of the model, there are many combinations of parameters (e.g. synaptic

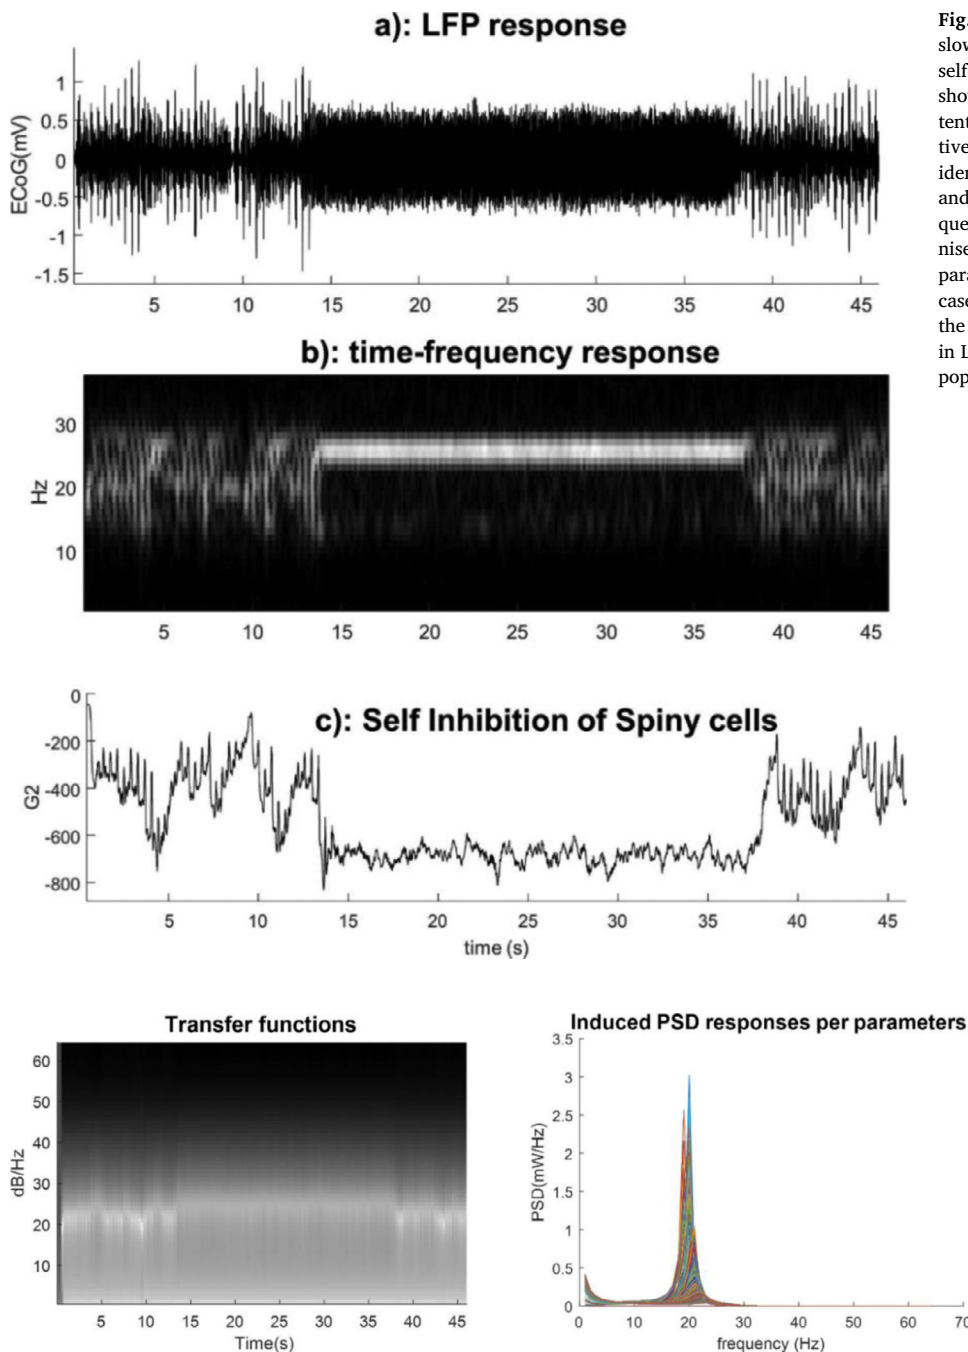


Fig. 3. Forward simulation of the CMC model with slow dynamics. a) LFP response of the model where self-inhibition of spiny and inhibitory populations showed slow changes that depend on postsynaptic potentials generated by spiny and inhibitory cells, respectively. Clearly, there are two distinct dynamics can be identified in the time domain - spiky normal activity and synchronised activity in the middle. b) Time frequency representation shows the frequency of synchronised activity in the beta range. c) Slow fluctuation of parameters before and after phase transitions- in this case there is a clear separation between the value of the parameter that produced the two types of activities in LFP. In part c, G_2 is the self-inhibition of excitatory population.

Fig. 4. Forward simulation of the transfer function CMC with slow dynamics. The figure on the left-hand side shows the time frequency response of the transfer function given the trajectory of slow parameters in Fig. 2 panel c. The right-hand side of the figure shows the power spectral density of the simulated transfer function for each value of slowly varying parameters in Fig. 3 panel c (mapped into the time frequency domain on the left figure).

time constants, connectivity parameters) that could provide an equally good fit to the data (Jansen and Rit, 1995) and some may have even better free energy scores than the model generating the data (Litvak et al., 2019; Friston et al., 2013). This may sound counterintuitive; however, recall that the free energy is the trade-off between the accuracy and complexity of the model, where complexity is the difference (KL-divergence) between the priors and posteriors. Slight changes from the prior expectations in a large number of parameters may offer a less complex explanation for the data than having just a few parameters with a large deviation from their priors. Therefore, a higher evidence may be afforded to a model that is simpler than the one used to generate the data.

In this example, we first define a prior over all model parameters and then test various models in which one parameter is allowed to fluctuate around its prior expectation over epochs. To identify the prior expectation, we estimated all model parameters from a normal segment of the

data in Fig. 7 (alternatively, one could use several normal segments and use Bayesian model averaging to define the requisite expectation).

After identifying the prior expectation for the parameters, we ran DCM for CSD for each epoch to quantify changes in self-inhibition associated with each population. In total there are four self-inhibition connections. In other words, the hypothesis that we evaluate here is that pathological disinhibition in neuronal populations induces paroxysmal transitions and our key question of interest is: which populations contribute to the genesis of seizures? Answering such a question could have a substantive impact on our understanding of epileptic seizures (as opposed to non-epileptic seizures), which in turn may assist in designing an effective treatment strategy to suppress or abate pathology.

Practically, we took the above parameter estimates (posterior mean and covariance) to a second level parametric empirical Bayes (PEB) analysis (Friston et al., 2016), with a between-epoch design matrix contain-

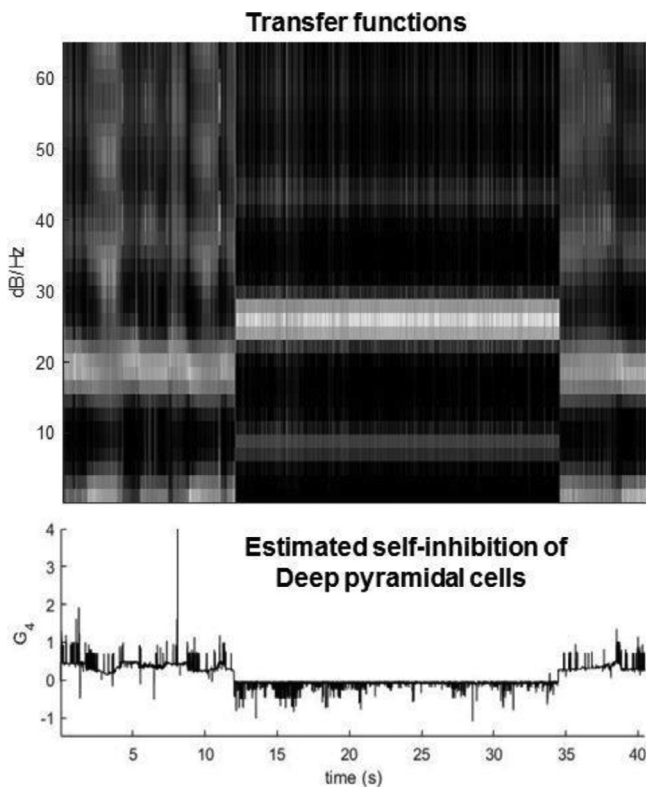


Fig. 5. Predicted time-frequency response of the estimated CMC over epochs. DCM for CSD is used to estimate parameters for each epoch. The ensuing estimates are then used to estimate the transfer functions of the CMC model, to generate predicted induced frequency responses. The lower panel shows the variation of parameter estimates in example 1. Note the quantitative difference between parameters during and after seizures. This figure shows the log scale parameter pertaining to intrinsic self-inhibition connection G_4 of deep pyramidal cells.

ing empirical priors based on the (binarized) spectral envelope of the data. These empirical priors tell us *when* particular synaptic parameters change; enabling model comparison to identify *which* particular combinations of parameters best explain the data. As it shown in Fig. 8b,

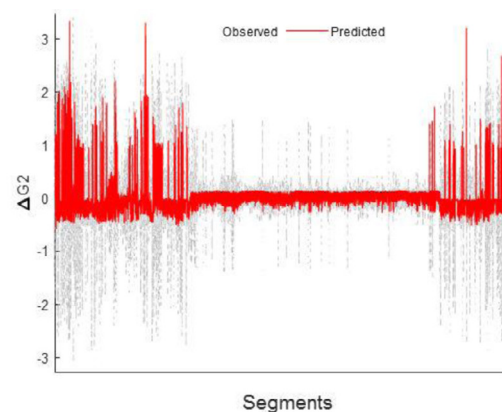
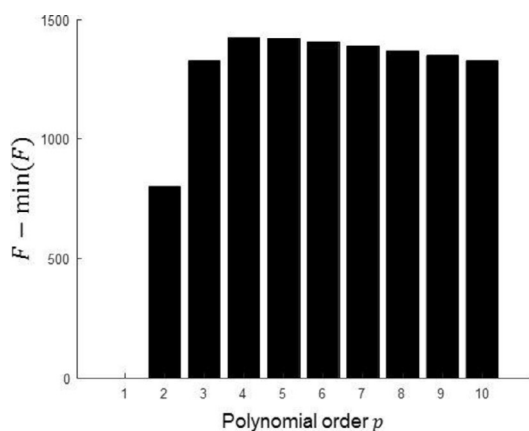


Fig. 6. Modelling the evolution of parameter (ΔG_2) with respect to empirical frequency content in the data. a) This panel shows relative free energy (with respect to the smallest model evidence) of the second level GLM. We estimated the model evidence associated with models of changes in synaptic parameters based on a linear-in-coefficients polynomial of the empirical frequency content of the ensuing neuronal responses – as captured by the data (Eq. 6). Here the optimal power of the polynomial is 4. b) Observed and predicted model of the parameter dynamic (ΔG_2) with polynomial order 4 in Eq. 6. In this and all subsequent figures, changes in the parameters are expressed in terms of log scaling. In other words, a value of 0 corresponds to a scaling by $\exp(0)=1$. For small changes, the log scaling is roughly equal to the proportional change (i.e., -0.1 is roughly a decrease of 10%). The free energy of a model with order 4 was 30674 and for order 5 was 30557 < 30674. By converting these log evidence approximations to the posterior probability of each model (which by Bayes rule under equal priors is a softmax function), we would select the model order of 4, which has posterior probability approaching 100%.

Bayesian model reduction suggests that changes in self-inhibition in inhibitory and deep pyramidal cells are the best explanation for the data. This is consistent with how the data were actually generated. Bayesian model reduction is used to eliminate redundant parameters by testing the evidence for models with and without a particular parameter and computing the Bayesian model average.

We then repeated the PEB analysis by replacing the binarized spectral envelope with frequency specific regressors, to characterise the relationship between the self-inhibition and spectral responses. This kind of analysis could disambiguate the contribution of distinct neuronal populations to paroxysmal brain activity (which can be targeted by different intervention mechanisms). We selected three frequencies as the regressors for the PEB design matrix (i.e., GLM). Specifically, we identified the two frequencies that predominated during the seizure period and normal activity. Then used PEB to characterise the contribution of parametric changes to spectral responses in the data (or vice versa). The result of this analysis is shown in

Fig. 9, and suggests that self-inhibition of the inhibitory population is likely to be responsible for the generation of 3 Hz oscillation (or vice versa) and both inhibitory and deep pyramidal self-connections are implicated in the generation of 8 Hz activity. Fig. 9 shows the posterior estimates of the second level (GLM) parameters (on top) and following Bayesian model reduction (below).

The inferred parameters of the GLM (coupling frequency-specific activity to self-inhibition parameters) are plotted in Fig. 8. As can be seen from these posterior estimates, the upward/downward changes in self-inhibition of the inhibitory and deep pyramidal cells are well captured. This frequency-specific analysis may provide valuable information that it would not otherwise be possible to extract from electrophysiological data. More importantly, capturing trends in parameter dynamics may be important for designing effective treatments of epilepsy as different intervention mechanisms may have an opposite effect on different ion currents (Blenkinsop et al., 2012; Nevado-Holgado et al., 2012). In the final section, we rehearse A-DCM, using empirical data to provide an illustrative example.

4. Worked example: animal model of epilepsy

As the final example in this paper, we apply A-DCM to a Picrotoxin animal model of seizures (Kätzel et al., 2014). In this animal model, the cause of paroxysmal activity can be related to pathological disruption of

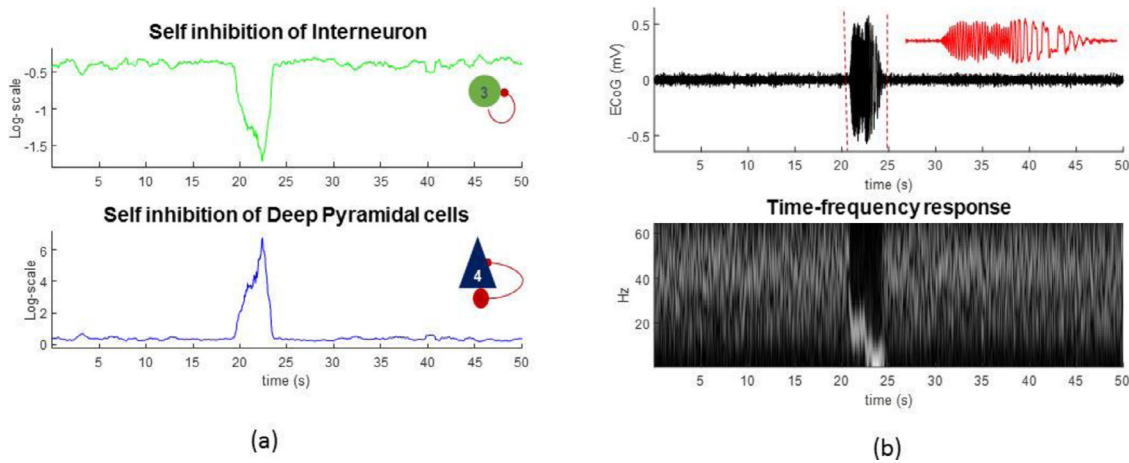


Fig. 7. Simulation of induced seizures. a) Slow dynamics of parameters of the CMC model. b) Simulated ECoG and its time frequency representations, which shows some features of spontaneous seizures, i.e., before the seizures we observe fast activity that is followed by pathological spikes around 5-8 Hz.

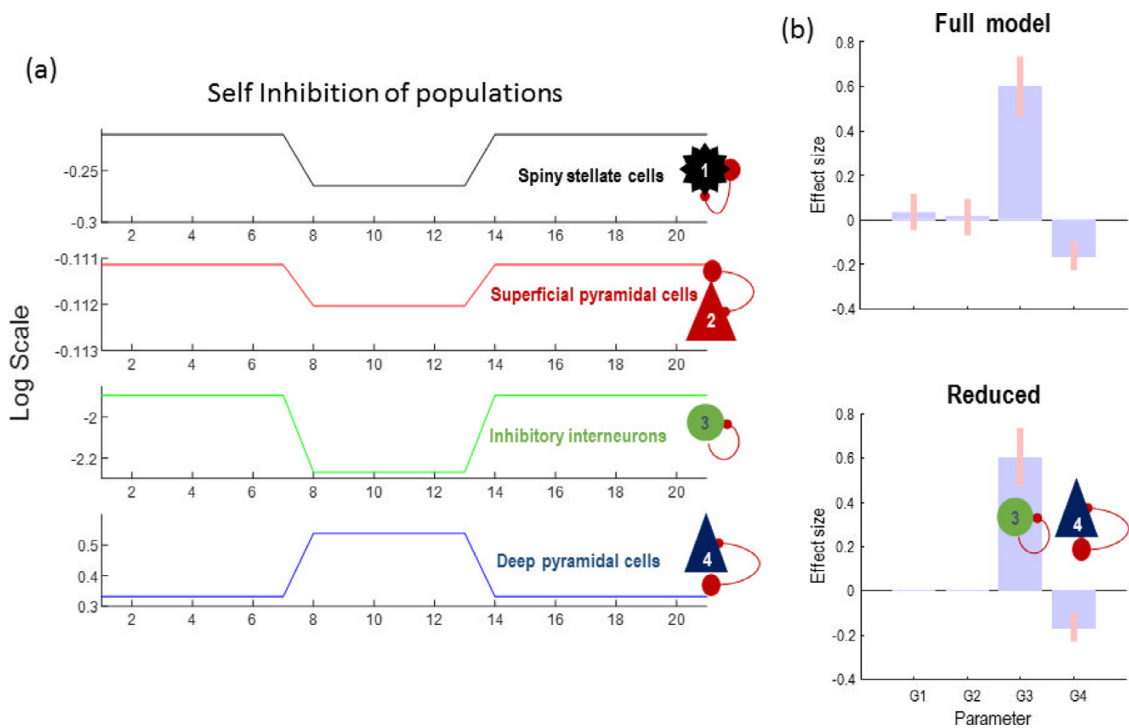


Fig. 8. Parametric empirical Bayes analysis of simulated data. a) Estimated slow dynamics of parameters in the CMC mode over 21 segments of data with average of 2 second duration. b) PEB analysis with a design matrix based on the spectral envelope of the data. The original PEB posterior estimates show on top (full model) and the estimates following Bayesian model reduction are shown below (reduced). The reduced model reveals the underlying causes of seizures in Fig. 7 and clearly capture is the trends in parameters. G_1 , G_2 , G_3 and G_4 are self-inhibition of excitatory, superficial pyramidal, inhibitory, and deep pyramidal cells, respectively. As in previous figures, the effect sizes relate to unitless log scaling parameters.

synaptic connections at the focus of drug injection, which is then spread to other regions (Kätzel et al., 2014; Wood, 2014).

The animal was implanted with a wireless ECoG device on its brain and a cannula was inserted into its primary visual cortex, in the right hemisphere. We then injected a chemo-convulsant drug (Picrotoxin, 300-400nl of 10mM) via the pre-implanted cannula. A few minutes after the injection, large amplitude spikes emerged (due to the hyper-synchronised response of the neuronal population at the site of intervention), followed by occurrence of seizures ~20-30 minutes after injection. Behavioural manifestation of seizures includes freezing, followed by head bobbing (20-30s), hunching, limb-kicking and occasionally rearing, falling over and wet dog shakes. Average dura-

tion of the seizures period in this animal model is around ~ 53s, with an inter seizure interval of approximately ~200 s. In this animal model, the activity of the brain recovers to the baseline activity 2-3 hours after the injection, which implies ~15-20 seizures over 2 hrs. A sample of ECoG data and its time frequency response are shown in Fig. 10.

At the initial phase of drug injection, seizures commenced focally and only the injection site was affected by the drug. Subsequently other brain regions were recruited by the seizures. We elected to study the underlying causes of the seizures in the initial phase of the pathology and investigate its underlying mechanisms to illustrate an application of A-DCM. See also Appendix B.

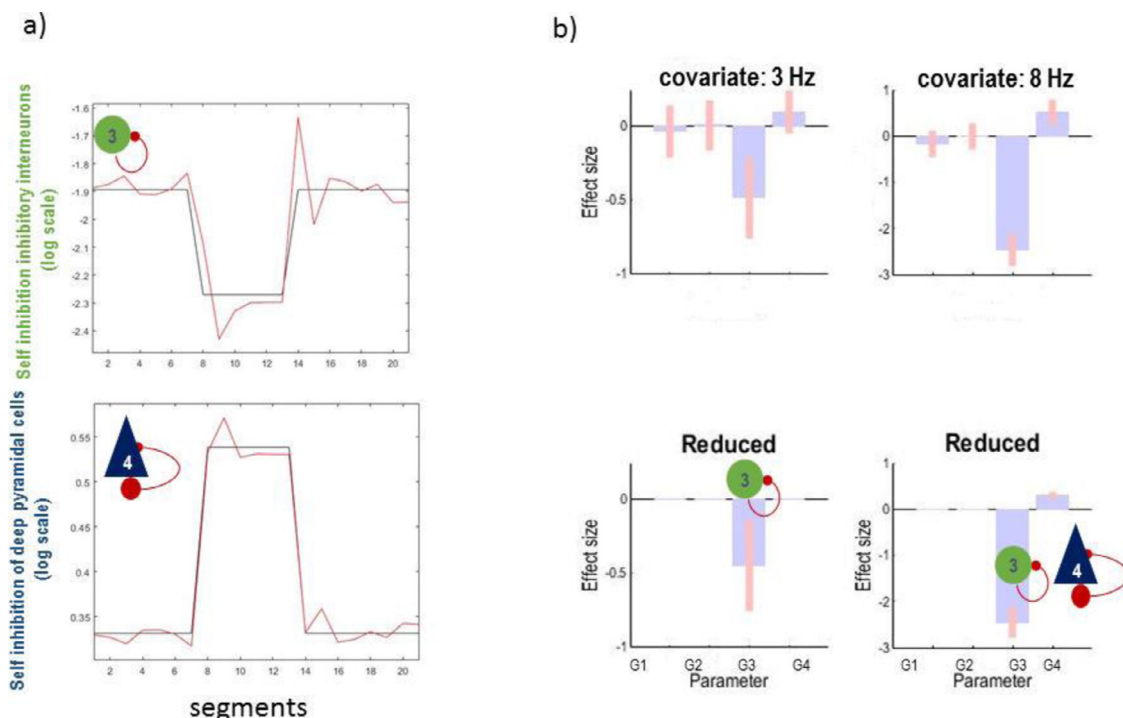


Fig. 9. Parametric empirical Bayes analysis based on the predominant frequency content of the data. a) changes in parameters (black) and predicted changes of G_3 and G_4 (red) using PEB with three frequency covariates in the GLM design matrix. b) second level GLM parameters (above) and following Bayesian model reduction (below). The BMR suggests that slow waves (3 Hz) are associated with inhibitory populations, whereas higher frequency pathological activity (8 Hz) may be induced by both deep pyramidal and inhibitory populations. G_1 , G_2 , G_3 and G_4 are self-inhibition of excitatory, superficial pyramidal, inhibitory, and deep pyramidal cells, respectively.




Fig. 10. Real ECoG data from an animal model with normal to large spikes and seizures. Lower panel shows a scaled time-frequency plot of the ECoG data.

We first fitted the CMC model to the normal activity to establish the prior expectation for model parameters. We then epoched the data each of which has 2 seconds duration and estimated key model parameters. We fixed some of the parameters (e.g., noise hyperparameters and sensor gain, which are not likely to vary in this animal model during

the experiment) and only allow the rate constant and self-inhibition of neuronal populations to change during and after seizures. After model inversion, we ran a PEB analysis with a binarized spectral envelope as a regressor (column vector with zero and one entries, where zero and one denotes normal and seizure epochs, respectively). The results are shown in Fig 11 and suggest that disinhibition in the inhibitory and deep pyramidal cells populations best explain the data.

Our results are in agreement with the clinical understanding of seizures in this animal model (Wood, 2014; Berglind et al., 2014). In detail, these seizures are thought to be due to disruption of interaction between inhibitory populations and other populations (Berglind et al., 2014; Rovainen, 1983). In our modelling results, seizures were explained by reduction of inhibitory self-connectivity (i.e., disinhibition). These results are informative because disruption in the self-regulation of the inhibitory populations engenders changes in postsynaptic interneuron potentials, and thereby affects activity of other populations. Our analysis also suggests that the deep layer is likely to contribute to seizures dynamics, which is plausible, since from an anatomical perspective the density of neurons in the deep layer of the cortex is greater than in the superficial layers. In this DCM, such mechanisms are reflected in the synaptic time constant of deep pyramidal populations that increased by over 50% in this example.

Finally, we repeated the above analysis using frequency-specific regressors to characterise the relationship between disinhibition in particular populations and their frequency specific correlates. Here, the regressors of the GLM were the predominant empirical frequencies in the data of 5, 19, and 40 Hz (identified using singular value decomposition of the time frequency data). The PEB results (Fig. 12) show that seizures (which are characterised by 3 to 8 Hz activity) are best explained by the reduction in effective membrane time constant of deep pyramidal cells, with interneurons contributing to 19 Hz activity. This is again consistent with physiological findings from this animal model of seizures.

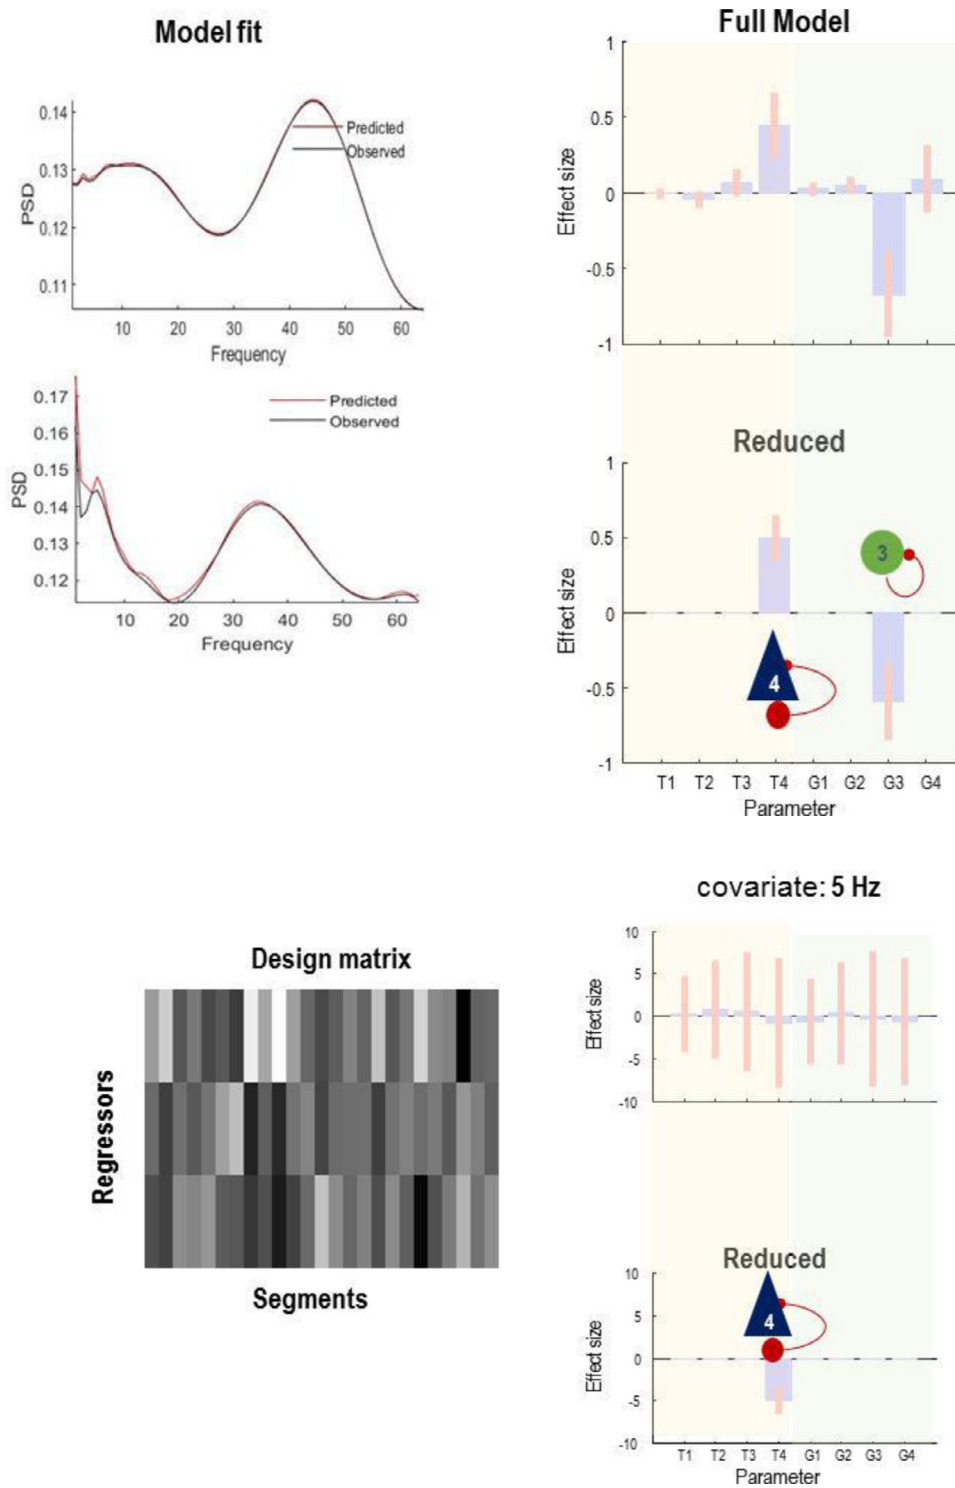


Fig. 11. Model fit plots that show predicted and observed signals using the DCM for CSD approach for example normal and pathological segments of the animal data. The right panel shows the results of parametric empirical Bayes (PEB) and Bayesian model reduction (BMR) of slowly changing parameters. The PEB model is a hierarchical model, with a general linear model (GLM) of the neural parameters at the between-epoch level. The design matrix of the GLM is simply the binarized spectral envelope encoding epochs with and without seizure activity. BMR suggests that disinhibition of deep pyramidal and inhibitory populations was likely to explain the seizure activity. G_1 , G_2 , G_3 and G_4 are self-inhibition parameters of excitatory, superficial pyramidal, inhibitory, and deep pyramidal cells, respectively. T_1 , T_2 , T_3 and T_4 are the time constants of excitatory, superficial pyramidal, inhibitory, and deep pyramidal cells, respectively.

Fig. 12. Parametric empirical Bayes and Bayesian model reduction modelling of slowly changing synaptic parameters. The regressors of the between-epoch GLM comprise the frequency content in the signal (left hand side of the figure). The right side of the figure shows the GLM parameters (top row) and the reduced parameters after Bayesian model reduction (bottom row). The model reduction with this design matrix shows that inhibition and deep pyramidal cells are likely to explain the seizures at the different frequencies. G_1 , G_2 , G_3 and G_4 are self-inhibition parameters of excitatory, superficial pyramidal, inhibitory, and deep pyramidal cells, respectively. T_1 , T_2 , T_3 and T_4 are the time constant of excitatory, superficial pyramidal, inhibitory, and deep pyramidal cells, respectively.

5. Discussion

This paper has introduced a adiabatic dynamic causal model (A-DCM) that enables one to compare competing hypotheses about the biological mechanisms that might underwrite phase transitions in electro-

physiological recordings. Crucially, A-DCM is formulated to elucidate the most likely causal relationships between synaptic efficacy and spectral activity in electrophysiological data, which is commonly used characterise and understand brain states in the healthy and pathological brain (Shaw et al., 2017). The use of an adiabatic approximation and

mean field theory furnishes an efficient way to infer the relationship between slow changes in synaptic parameters and their neuronal correlates in spectral data. The resulting model could, in principle, be used to infer how interventions modulate states of the diseased or healthy brain (Fleming et al., 2020; Liang et al., 2015; Moran, 2015). This would require construct validation experiments to test consistency and agreement between predictions of the model and effects of interventions on real brains e.g., (Mina et al., 2013).

We have motivated A-DCM using Synergetic theory and the Adiabatic approximation. As touched on in the introduction, A-DCM can potentially provide complementary information alongside the conventional slow-fast modelling approach. In A-DCM slow variables are modulated with respect to the frequency contents of fast states, the slow variable is equipped with a sort of memory (since the frequency domain representation of neuronal responses rests on the second order statistical moments of fast states). In this respect, one can establish a link between parametrised fluctuations in A-DCM and slow dynamics in slow-fast dynamical systems (Sanders et al., 2007). Akin to A-DCM idea, averaging of slow-fast dynamical systems allows separation of slow and fast dynamics, where the slow dynamics are an integral of fast states (under the assumption that fast states are ergodic for each value of the slow states). We also note conceptual links between A-DCM and linear response theory (LRT) in physics (Ruelle, 2009). Specifically, LRT implies that if an ergodic system – e.g., neural mass or mean field model (Marreiros et al., 2009) – is left without any perturbation, it eventually reaches its equilibrium (Lucarini, 2008). If, however, some parameters are slowly changing, LRT assures the existence of a new equilibrium for the system that can be attained instantaneously. In addition, LTR suggests that a process with slowly varying parameters can be reversible (i.e., it can attain its initial equilibrium). Crucially and more closely related to A-DCM, LTR establishes a link between time domain features of dynamical systems and its equilibrium, thereby allowing parametric/functional expression of the system's response in the time/frequency domain (Reick, 2002; Lucarini, 2008; Lucarini et al., 2007; Dykman et al., 1998).

We have introduced the notion of circular causality in a slightly 'tongue-in-cheek way'. In a dynamical setting, circular causality depends upon bottom-up and top-down causation among the latent states generating data. This entails a generative model in which fast neuronal states cause slow, activity-dependent changes in synaptic parameters; while, at the same time, the synaptic parameters enslave the fast, neuronal states. In our adiabatic DCM estimation, the coupling between fast and slow dynamics is modelled by generating fast data features using slow data features as priors on the latent states. As noted above, this means the data are used twice: as both cause and consequence of fluctuations in hidden (neuronal) states and parameters. However, the data features in question are not the same. The data being predicted summarise fast neuronal activity in terms of its spectral content within an epoch. Conversely, the empirical priors are furnished by slower changes in spectral content from epoch to epoch. This separation of temporal features characterises this particular DCM and sets it apart from previous models, upon which the current adiabatic DCM is based. In summary, our focus was on establishing casual links between slow and fast states, both in terms of forward simulation and estimation procedures. To do this, we used models that were equipped with slowly varying parameters, which cause – and are caused by – changes in the spectral contents of fast states. This builds upon the use of DCM to characterise dynamic effective connectivity (Park et al., 2017; Van de Steen et al., 2019). For example, (Park et al., 2018), used a principal component analysis of the first level connectivity parameters to generate empirical (between-epoch) priors for time-varying (within-epoch) effective connectivity. In the present setting we used the spectral features of empirical responses to furnish empirical priors.

In this work, we offered two forward simulations to illustrate the different kinds of phase transition that one might infer using A-DCM. In the first simulation, phase transitions in brain activity were explained when parameters moved from one region of parameter space to another, i.e.

bifurcations (Breakspear et al., 2006). In the second simulation, we show that underlying causes of paroxysmal transitions can be understood as a hysteresis phenomenon (Iasemidis, 2003; Voss et al., 2012), where in one direction of parameter variations, the model produces gamma activity while in reciprocal direction, the model generates pathological spike-wave discharge activity.

In terms of estimating parameters, in the first face validity analysis, we illustrated how one can establish links between fluctuation of parameters and spectral responses in the data. This perhaps would be most useful for tracking dynamics of parameters with respect to changes in spectral responses in real data. The potential application of such a generative model would be tracking, where we would be interested in the modulation of brain activity with respect to changes of transmembrane currents, which may be subject to alteration by interventions. In the second face validity study, we recovered parameter trends that induced spontaneous seizures. In addition, we showed that it is feasible to use Bayesian model reduction to evince underlying biological mechanisms that explained transition dynamics. In principle, one could ask whether seizure dynamics are noise-driven, as opposed to itinerant dynamics (e.g., heteroclinic cycles in parameter space). Because the generative model of DCM for CSD is equipped with (a model of) random neuronal fluctuations, one can fit separate DCM for CSD models with neuronal fluctuations that are fixed or time varying across epochs, and compare their evidence using Bayesian model comparison. Clearly, this would rest on the assumption that seizure activity was accompanied by changes in the level of random fluctuations; enabling neuronal dynamics to explore other basins of attraction. Furthermore, the linearisation – inherent in estimating epoch-specific parameters – precludes an explicit modelling of deterministic multi-stability. At first glance, this may appear a limitation of linearised models. However, stochastic chaos is distinct from multi-stability in deterministic systems and may be approximated more easily with local linear assumptions. For example, bi-stability in a deterministic system with two basins of attractions will, with an appropriate level of random fluctuations, present as a single pullback attractor (Crauel and Flandoli, 1994). Please see (Moran et al., 2011) for a discussion of these and related issues.

The temporal resolution of the inferred parameters depends upon the number of epochs and the overlap between epochs. In this paper, the overlap between epochs was based on forward simulation of the model. For instance, as shown in Fig. 3, the parameters were varied very quickly from one epoch to the next (which in turn induced spikes in pre-seizure activity), whereas in Fig. 7, the parameters changed slowly and smoothly (prior to seizure onset). Therefore, we used overlapping epochs to infer parameters that reproduced the data in Fig. 3, whereas to estimate model parameters for the data in Fig. 6, a non-overlapped approach was employed. Having said this, it would be possible to use a pragmatic approach to segmenting data into sub (quasi)-stationary segments using change point detection algorithms (a function of this kind is now implemented and available in Matlab) to define boundaries of epochs for the data with pathological transitions—and one may overlap between them (to assure proper estimation of PSDs).

In the worked example, we used the CMC model to model the underlying causes of drug-induced seizure activity. Here our results were in line with known pathophysiology. In this example, the underlying causes of seizures were attributed to disruption of inhibition. Our modelling approach suggests that the inhibitory population is largely responsible for the generation of 3 Hz waveforms. One could ask whether other – potentially better – explanations could be found for the observed data. Here, we limited the space of possible hypotheses, by using an animal model, where seizure aetiology was known to have its source at the injection site. Our results were consistent with what is known about this animal model. Nevertheless, better explanations for these data might be found, either through different mixtures of connectivity parameters in the same kind of model, or by selecting a model with a different form. For example, one might try conductance-based models, which quantify changes of ion dynamics. The confidence we can place in these differ-

ent explanations for the data is based on the posterior probability of the accompanying models. Predictive validity can then be assessed by asking whether artificial lesions or drugs administered to the DCM generate similar changes to data features seen empirically. While our results clearly need to be further reproduced and validated; they speak to a promising application of A-DCM. As measured above, we only considered the initial phase of seizures in this animal model, where underlying causes of the seizures can be linked with the local effects of the drug. More interesting questions might be investigated using A-DCM in the second phase of seizures in this animal model, where other brain regions are contributing to the pathology.

This paper has focused on a single neural mass model and associated transfer functions. Clearly, one question that arises is: can this approach be scaled up to large networks or graphs? DCM for CSD has been applied to distributed brain networks in many applications e.g., (Papadopoulou et al., 2017; Rosch et al., 2018a; Rosch et al., 2018b). Due to the computational efficiency of variational approaches, there are no special constraints on the size of the network. However, a deeper issue here is the inherent complexity of the models needed to explain multichannel or multielectrode recordings. In other words, there may be an upper bound on the number of nodes – and their connections – that can be inferred on the basis of any given data. In principle, this bound would be addressed using Bayesian model comparison. In other words, there will be an optimum size of the network for any given set of data that would maximise Bayesian model evidence or its variational free energy bound.

A general issue with dynamic causal models is the potential for local minima during model inversion; especially when dealing with expressive models with large numbers of parameters. However, perhaps counterintuitively, more expressive models can elude local minima, because the parameter space has more dimensions or directions from which the search can escape. This means that normal practice would be to start with an over-parameterised model and then use Bayesian model reduction to eliminate redundant parameters. Bayesian model reduction in this setting is just an instance of Bayesian model comparison based upon model evidence. Crucially, the model is defined uniquely in terms of priors. This means that one can score the effect of changing the priors in terms of its effect on the model evidence—and thereby quantify the role of priors in explaining the data at hand.

A key theme in this modelling endeavour is the circular causality between connection strengths or synaptic efficacy and the neuronal activity these connections support. There is a large literature on models of synaptic plasticity (Demšar and Forsyth, 2020) and associative plasticity (Humeau et al., 2003; Kujirai et al., 2006); namely, the relationship between pre- and post-synaptic activity that can be linked to slow variation of synaptic parameters. One could also consider activity-dependent plasticity (Rebola et al., 2010; Isomura and Friston, 2019) as a function of the complex cross spectra or cross covariance functions summarising neuronal dynamics. In turn, this means that one could appeal to the notion of spike timing-dependent plasticity (STDP), to model changes in effective connectivity in terms of the cross-covariance function between a source and target population (which is well formulated in the spectral domain). The utility of expressing plasticity (i.e., the dynamics of slow parameters) in terms of covariance functions is that there is an equivalent frequency space representation; thereby accommodating the representation of dynamics in terms of complex cross spectra in the generative (dynamic causal) model. The only unknown hyperparameters in this instance are the coefficients scaling the amplitude and width of the functions that lead to increases and decreases in synaptic efficacy. Furthermore, there are empirical constraints on these functions. For example, for STDP, one would normally use the first derivative of a Gaussian function, with a dispersion of about 50 ms. With this formulation of adiabatic dynamics, we automatically account for ideas like spike rate adaptation (Peron and Gabbiani, 2009), spike timing dependent plasticity (Caporale and Dan, 2008; Dan and Poo, 2006) and the transmission delays inherent in the dynamic causal modelling of cross

spectra (Friston et al., 2012). In the future, we hope to extend A-DCM to address the delicate interplay between synaptic plasticity and spectral fluctuations. The second level PEB design matrix (which is used to constrain model parameters based on the spectral content of the data) can also be augmented with regressors that encode information about fluctuations in other high-order statistics (e.g., Lyapunov exponents, Hausdorff dimension, et cetera). Leveraging such prior information may help to capture the underlying dynamics of seizure initiation, termination and related phase transitions.

Credit authorship contribution statement

Amirhossein Jafarian: Conceptualization, Methodology, Software, Validation, Formal analysis, Writing – original draft, Data curation, Visualization. **Peter Zeidman:** Conceptualization, Methodology, Writing – review & editing, Supervision, Project administration, Investigation. **Rob. C Wykes:** Conceptualization, Writing – review & editing, Supervision, Project administration, Investigation. **Matthew Walker:** Writing – review & editing, Supervision, Investigation, Conceptualization. **Karl J. Friston:** Conceptualization, Methodology, Writing – review & editing, Supervision, Funding acquisition.

Acknowledgments

The Wellcome Centre for Human Neuroimaging is supported by core funding from Wellcome [203147/Z/16/Z]. RW holds a Senior Research Fellowship funded by the Worshipful Company of Pewterers.

Data and code availability statement

The code associated with this paper will be available in the next release of the SPM12 software (<https://www.fil.ion.ucl.ac.uk/spm/>).

Appendix A

A technical aspect of this generative model is the form of the likelihood for the complex cross spectra. With ideal estimators, one could assume that these spectral data features had a Wishart distribution, with one degree of freedom for each frequency. However, we can assume that the cross spectra constitute the average of estimates, with consistent and asymptotically normal estimates. In this setting, the variance of the difference between the predicted and observed spectral estimates is equal to the cross spectral density times the (effective) degrees of freedom (h), which we treat as an unknown parameter. Specifically, the precision is given by the asymptotic results where, the scaled difference between the sample spectral density (g) and the predicted density (G) (Camba-Mendez and Kapetanios, 2005):

$$e = \text{vec}(g - G) \quad (\text{A.1})$$

is asymptotically complex normal and the covariance between $e(i, j)$ and $e(u, v)$ is given by Q/h and:

$$Q = G(i, u) * G(j, v); \quad h = 2 \times m + 1$$

Here m represents the number of averages from a very long time series. The inverse of the covariance is thus a scaled precision, where the hyperparameter (h) plays the role of the degrees of freedom (e.g., the number of averages comprising the estimate). We use the sample spectral density to create a frequency specific precision matrix for the vectorised spectral densities, under the assumption that the form of this sample spectral density resembles the predicted spectral density (which will become increasingly plausible with convergence).

Appendix B

In the animal model of seizures in this paper, paroxysmal activity may be induced as a result of the disruption of the interaction between

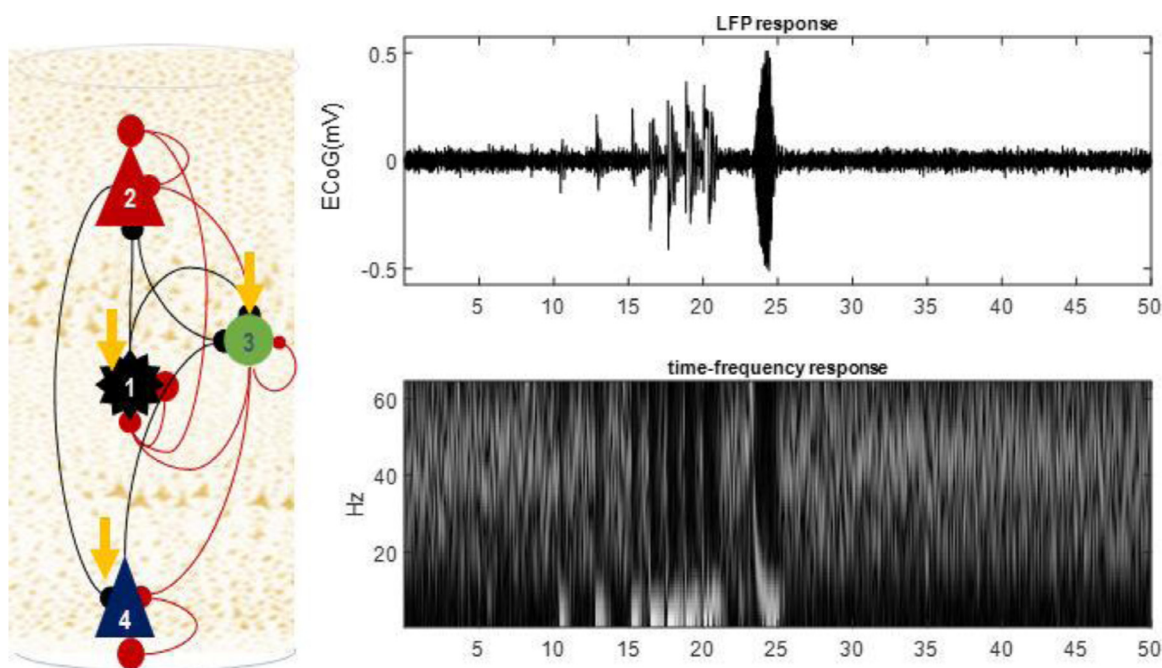


Fig. 13. Forward simulation of the CMC model with time varying parameters. Here the rate constants and synaptic efficacy of inhibitory, spiny stellate and deep pyramidal cells are changed from their prior expectation to generate spike-like activity and recurrent of seizures, as shown in LFP response of model.

inhibitory interneurons with other populations. Such disruption can be modelled by changing the balance between synaptic efficacies and gains of excitatory and inhibitory populations (Wendling et al., 2002). For example, since the drug was delivered to the deep layer of cortex, one may alter the rate constant/self-inhibition of inhibitory, spiny stellate excitatory cells and deep pyramidal cells in the CMC model to replicate both spike discharges and recurrent of seizures. One example of such a forward simulation is shown in Fig. 13.

One should note that, this is only one way to generate pathophysiology and there may be many other plausible mechanisms that one could implement using the CMC model to generate paroxysmal transitions of the data. As explained in the paper, the most compelling questions – for epileptologists – are which synaptic parameters are likely to explain the underlying causes of pathognomonic data?

Appendix C

Many mathematical treatments of paroxysmal activity are cast in terms of bifurcations and phase transitions – borrowing techniques from dynamical systems theory (e.g., (Saggio et al., 2020)). The linearised models adopted by the DCM for CSD approach preclude an explicit formulation in terms of dynamical instability or multi-stability. Furthermore, the underlying generative model is predicated on ensemble dynamics that may or may not evince stochastic chaos (as opposed to deterministic chaos). Having said this, the notion of bifurcation still plays an important role in understanding the way that DCM generates data. This is meant in the sense of transcritical bifurcations (normally hyperbolic fixed points), where movement in parameter space causes a threshold crossing of one or more eigenvalues of the system's Jacobian. These threshold crossings are usually suppressed in DCM – and modelled as critical slowing by allowing for eigenvalues that approach zero from below but never reach zero (i.e., the generative model assumes, a priori, self-organised criticality but precludes trans-critical bifurcations). The ensuing dynamics therefore approach a transcritical bifurcation producing similar spectral data features, while retaining dynamical stability. Put more simply, DCM can reproduce the spectral features that are usually associated with bifurcations in deterministic systems but does so by using a local linear approximation that is always dynamically stable.

This means that seizure onsets and offsets can be modelled by movement in parameter space that approximates, to first order, the kind of phase transitions seen in deterministic formulations.

At the onset/offset of seizures dynamics, rapid changes in brain states—from normal to pathological (and vice versa) makes it challenging to track very fast changes in model parameters (i.e., synaptic efficacy) that could underwrite phase transitions. However, one could smooth spectral transitions. Specifically, to model the onset (offset) of paroxysmal transition, one could generate so called geodesic paths between power spectral densities (PSDs) of normal and seizures activities to identify changes in model parameters that induce transitions: let the PSDs of normal and seizure activity be denoted by $P_1(\omega)$, $P_2(\omega)$ respectively; then parameterised geodesic paths between these spectral densities can be generated via $P_\tau(\omega) = P_1^\tau(\omega)P_2^{1-\tau}(\omega)$, where $\tau \in (0, 1)$ (Georgiou, 2007). By selecting different values for τ , one could generate spectral data features that connect the normal and seizure activity in the frequency domain. Parameters that model these spectral features should trace a path in parameter space that can explain the onset (offset) of seizure activity.

Appendix D

Table 1D This appendix summarises the canonical microcircuit (CMC) (Bastos et al., 2015; Bastos et al., 2012) that we employed in this paper, and the text is adapted from (Jafarian et al., 2020). The CMC model comprises four neuronal populations: namely, spiny stellate cells in the granular layer (ss), superficial pyramidal cells in the supragranular layer (sp), inhibitory interneurons distributed in all layers of the cortex (ii) and deep pyramidal cells in the infragranular layers (dp). In this model, two conversion operators govern the dynamics of each neuronal population (Jansen and Rit 1995). The first operator converts the mean pre-synaptic firing rate m to the mean postsynaptic membrane potential V as follows (Freeman, 1975):

$$V = h \otimes m \quad (1d)$$

Where \otimes denotes the linear convolution operator and h is the impulse response function (synaptic kernel) with synaptic rate constant T :

Table 1D
Default values for the parameters of the CMC model.

Description	Value [ss sp ii dp]
T	Postsynaptic rate constant of the i -th neuronal population.
T	$T = [256, 128, 16, 32]$
$g_{i \rightarrow k}$	Intrinsic connectivity between populations i and k (self-inhibition in case of $i=k$ which would be denoted by $G_{k \rightarrow k}$).
	$g = [2 \ 1 \ 1 \ 1] * 512$

$$h(t) = \begin{cases} \frac{t}{T} e^{-\frac{t}{T}}, & x \geq 0 \\ 0, & x < 0 \end{cases} \quad (2d)$$

The second operator then transforms the postsynaptic membrane potential into a firing rate, which forms the input to the next connected neural population (Jirsa and Haken, 1997):

$$\sigma(V) = \frac{1}{1 + \exp(-V)} - \frac{1}{2} \quad (3d)$$

In effect, the dynamics of postsynaptic potentials in the population i , V_i , obey the following second order differential equations as follows:

$$\left(1 + \frac{1}{T_i} \frac{d}{dt}\right)^2 V_i(t) = s(V_i, V_j) + \omega \quad (4d)$$

where the intrinsic presynaptic excitations from population j is denoted by V_j ; and the function s is defined as follows (Friston et al., 2019):

$$s(V_i, V_j) = \begin{cases} -G_{ss \rightarrow ss} \sigma(V_{ss}) - g_{sp \rightarrow ss} \sigma(V_{sp}) - g_{ii \rightarrow ss} \sigma(V_{ii}) + \omega & \text{if } i = ss \\ -G_{sp \rightarrow sp} \sigma(V_{sp}) + a_{ss \rightarrow sp} \sigma(V_{ss}) - g_{ii \rightarrow sp} \sigma(V_{ii}) & \text{if } i = sp \\ -G_{ii \rightarrow ii} \sigma(V_{ii}) - g_{dp \rightarrow ii} \sigma(V_{dp}) + g_{ss \rightarrow ii} \sigma(V_{ss}) + g_{sp \rightarrow ii} \sigma(V_{sp}) & \text{if } i = ii \\ -G_{dp \rightarrow dp} \sigma(V_{dp}) - g_{ii \rightarrow dp} \sigma(V_{ii}) + g_{sp \rightarrow dp} \sigma(V_{sp}) & \text{if } i = dp \end{cases} \quad (5d)$$

The laminar specificity of the extrinsic and intrinsic connections in Eq. (3) are specified by placing prior constraints on the intrinsic (within-region) connectivity parameters g , $G_{* \rightarrow *}$. The random neuronal fluctuations ω that drive the model have zero mean and small variance equal to $\exp(-32)$ in the simulations of this paper. The initial states of the model were zero. The initialisation of variational inference uses the prior expectation of the parameters. These priors – for the neural mass models at hand – have been chosen based on years of experience with this canonical microcircuit model and were originally based upon neurophysiology. Please see (Friston et al., 2019; Aukstulewicz and Friston, 2015; Bastos et al., 2015) for quantitative details.

References

Aukstulewicz, R., Friston, K., 2015. Attentional enhancement of auditory mismatch responses: a DCM/MEG study. *Cereb. Cortex* 25, 4273–4283.

Baier, G., Goodfellow, M., Taylor, P.N., Wang, Y., Garry, D.J., 2012. The importance of modeling epileptic seizure dynamics as spatio-temporal patterns. *Front. Physiology* 3, 281.

Basar, E., Flohr, H., Haken, H., Mandell, A., 2012. *Synergetics of the Brain: Proceedings of the International Symposium on Synergetics at Schloß Elmau, Bavaria, May 2–7, 1983*. Springer Science & Business Media.

Bastos, A.M., Litvak, V., Moran, R., Bosman, C.A., Fries, P., Friston, K.J., 2015. A DCM study of spectral asymmetries in feedforward and feedback connections between visual areas V1 and V4 in the monkey. *Neuroimage* 108, 460–475.

Bastos, A.M., Urey, W.M., Adams, R.A., Mangun, G.R., Fries, P., Friston, K.J., 2012. Canonical microcircuits for predictive coding. *Neuron* 76, 695–711.

Beal, M.J., 2003. Variational algorithms for approximate Bayesian inference. UCL (University College London).

Berglind, F., Ledri, M., Sørensen, A.T., Nikitidou, L., Melis, M., Bielefeld, P., Kirik, D., Deisseroth, K., Andersson, M., Kokaia, M., 2014. Optogenetic inhibition of chemically induced hypersynchronized bursting in mice. *Neurobiol. Dis.* 65, 133–141.

Bishop, C.M., 2006. *Pattern recognition and machine learning*. Springer.

Blenkinsop, A., Valentin, A., Richardson, M.P., Terry, J.R., 2012. The dynamic evolution of focal-onset epilepsies—combining theoretical and clinical observations. *Eur. J. Neurosci.* 36, 2188–2200.

Breakspear, M., Roberts, J., Terry, J.R., Rodrigues, S., Mahant, N., Robinson, P., 2006. A unifying explanation of primary generalized seizures through nonlinear brain modeling and bifurcation analysis. *Cereb. Cortex* 16, 1296–1313.

Camba-Mendez, G., Kapetanios, G., 2005. Estimating the rank of the spectral density matrix. *J. Time Ser. Anal.* 26, 37–48.

Caporale, N., Dan, Y., 2008. Spike timing-dependent plasticity: a Hebbian learning rule. *Annu. Rev. Neurosci.* 31, 25–46.

Carmichael, D.W., Vulliemmoz, S., Murta, T., Chaudhary, U., Perani, S., Rodionov, R., Rosa, M., Friston, K., Lemieux, L., 2017. Measurement of the mapping between intracranial EEG and fMRI recordings in the human brain. *bioRxiv*, 237198.

Chen, C., Kiebel, S.J., Friston, K.J., 2008. Dynamic causal modelling of induced responses. *Neuroimage* 41, 1293–1312.

Coombes, S., Bressloff, P.C., 2005. *Bursting: the genesis of rhythm in the nervous system*. World Scientific.

Crauel, H., Flandoli, F., 1994. Attractors for random dynamical systems. *Probability Theory Relat. Fields* 100, 365–393.

Da Silva, F.L., Hoeks, A., Smits, H., Zetterberg, L., 1974. Model of brain rhythmic activity. *Kybernetik* 15, 27–37.

Dan, Y., Poo, M.-M., 2006. Spike timing-dependent plasticity: from synapse to perception. *Physiol. Rev.* 86, 1033–1048.

David, O., Kilner, J.M., Friston, K.J., 2006. Mechanisms of evoked and induced responses in MEG/EEG. *Neuroimage* 31, 1580–1591.

Demšar, J., Forsyth, R., 2020. Synaptic Scaling Improves the Stability of Neural Mass Models Capable of Simulating Brain Plasticity. *Neural Comput.* 32, 424–446.

Du, M., Li, J., Wang, R., Wu, Y., 2016. The influence of potassium concentration on epileptic seizures in a coupled neuronal model in the hippocampus. *Cognitive Neurodyn.* 10, 405–414.

Dykman, M., Maloney, C., Smelyanskiy, V., Silverstein, M., 1998. Fluctuational phase-flip transitions in parametrically driven oscillators. *Phys. Rev. E* 57, 5202.

El Houssaini, K., Bernard, C., Jirsa, V.K., 2020. The Epileptor model: a systematic mathematical analysis linked to the dynamics of seizures, refractory status epilepticus and depolarization block. *Eneuro*.

Fleming, J.E., Dunn, E., Lowery, M.M., 2020. Simulation of closed-loop deep brain stimulation control schemes for suppression of pathological beta oscillations in Parkinson's disease. *Front. Neurosci.* 14, 166.

Freeman, W.J., 1975. *Mass action in the nervous system*. Citeseer.

Friston, K., Daunizeau, J., Stephan, K.E., 2013. Model selection and gobbledygook: response to Lohmann et. *Neuroimage* 75, 275–278.

Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., Penny, W., 2007. Variational free energy and the Laplace approximation. *Neuroimage* 34, 220–234.

Friston, K., Parr, T., Zeidman, P., 2018. Bayesian model reduction. *arXiv preprint arXiv:1805.07092*.

Friston, K., Penny, W., 2011. Post hoc Bayesian model selection. *Neuroimage* 56, 2089–2099.

Friston, K., Zeidman, P., Litvak, V., 2015. Empirical Bayes for DCM: a group inversion scheme. *Front. Syst. Neurosci.* 9, 164.

Friston, K.J., 2011. Functional and effective connectivity: a review. *Brain Connectivity* 1, 13–36.

Friston, K.J., 2014. On the modelling of seizure dynamics. *Brain* 137, 2110–2113.

Friston, K.J., Bastos, A., Litvak, V., Stephan, K.E., Fries, P., Moran, R.J., 2012. DCM for complex-valued data: cross-spectra, coherence and phase-delays. *Neuroimage* 59, 439–455.

Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. *Neuroimage* 19, 1273–1302.

Friston, K.J., Li, B., Daunizeau, J., Stephan, K.E., 2011. Network discovery with DCM. *Neuroimage* 56, 1202–1221.

Friston, K.J., Litvak, V., Oswal, A., Razi, A., Stephan, K.E., Van Wijk, B.C., Ziegler, G., Zeidman, P., 2016. Bayesian model reduction and empirical Bayes for group (DCM) studies. *Neuroimage* 128, 413–431.

Friston, K.J., Preller, K.H., Mathys, C., Cagnan, H., Heinzle, J., Razi, A., Zeidman, P., 2019. Dynamic causal modelling revisited. *Neuroimage* 199, 730–744.

Friston, K.J., Trujillo-Barreto, N., Daunizeau, J., 2008. DEM: a variational treatment of dynamic systems. *Neuroimage* 41, 849–885.

Fritschy, J.-M., 2008. Epilepsy, E/I balance and GABA_A receptor plasticity. *Front. Mol. Neurosci.* 1, 5.

Fung, P., Robinson, P., 2014. Neural field theory of synaptic metaplasticity with applications to theta burst stimulation. *J. Theor. Biol.* 340, 164–176.

Gavish, M., Donoho, D.L., 2014. The optimal hard threshold for singular values is $\sqrt{2}$. *IEEE Trans. Inf. Theory* 60, 5040–5053.

Georgiou, T.T., 2007. An intrinsic metric for power spectral density functions. *IEEE Signal Process Lett.* 14, 561–563.

Grasse, D.W., Karunakaran, S., Moxon, K.A., 2013. Neuronal synchrony and the transition to spontaneous seizures. *Exp. Neurol.* 248, 72–84.

Grenier, F., Timofeev, I., Steriade, M., 2003. Neocortical very fast oscillations (ripples, 80–200 Hz) during seizures: intracellular correlates. *J. Neurophysiol.* 89, 841–852.

Haken, H., 1977. *Synergetics*. Phys. Bull. 28, 412.

Hashemi, M., Hutt, A., 2016. A thalamocortical feedback model to explain EEG during anesthesia. *Selforganization in Complex Systems: The Past, Present, and Future of Synergetics*. Springer.

Hashemi, M., Vattikonda, A., Sip, V., Guye, M., Bartolomei, F., Woodman, M., Jirsa, V., 2020. 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*, 116839.

- Humeau, Y., Shaban, H., Bissière, S., Lüthi, A., 2003. Presynaptic induction of heterosynaptic associative plasticity in the mammalian brain. *Nature* 426, 841–845.
- Iasemidis, L.D., 2003. Epileptic seizure prediction and control. *IEEE Trans. Biomed. Eng.* 50, 549–558.
- Isumura, T., Friston, K., 2019. Reverse engineering neural networks to characterise their cost functions. *bioRxiv*, 654467.
- Jafarian, A., Freestone, D.R., Nešić, D., Grayden, D.B., 2019a. Identification of A Neural Mass Model of Burst Suppression. In: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, pp. 2905–2908.
- Jafarian, A., Freestone, D.R., Nešić, D., Grayden, D.B., 2019b. Slow-Fast Duffing Neural Mass Model. In: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, pp. 142–145.
- Jafarian, A., Litvak, V., Cagnan, H., Friston, K.J., Zeidman, P., 2020. Comparing dynamic causal models of neurovascular coupling with fMRI and EEG/MEG. *Neuroimage*, 116734.
- Jafarian, A., Zeidman, P., Litvak, V., Friston, K., 2019c. Structure learning in coupled dynamical systems and dynamic causal modelling. *Philos. Trans. R. Soc. A* 377, 20190048.
- Jansen, B.H., Rit, V.G., 1995. Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biol. Cybern.* 73, 357–366.
- Jansen, R., Timmerman, J., Loos, M., Spijker, S., Van Ooyen, A., Brussaard, A.B., Mansvelder, H.D., Smit, A.B., De Gunst, M., Linkenkaer-Hansen, K., 2011. Novel candidate genes associated with hippocampal oscillations. *PLoS One* 6, e26586.
- Jirsa, V.K., Friedrich, R., Haken, H., Kelso, J.S., 1994. A theoretical model of phase transitions in the human brain. *Biol. Cybern.* 71, 27–35.
- Jirsa, V.K., Haken, H., 1997. A derivation of a macroscopic field theory of the brain from the quasi-microscopic neural dynamics. *Physica D* 99, 503–526.
- Jirsa, V.K., Stacey, W.C., Quilichini, P.P., Ivanov, A.I., Bernard, C., 2014. On the nature of seizure dynamics. *Brain* 137, 2210–2230.
- Kass, R.E., Raftery, A.E., 1995. Bayes factors. *J. Am. Statist. Assoc.* 90, 773–795.
- Kätzel, D., Nicholson, E., Schorge, S., Walker, M.C., Kullmann, D.M., 2014. Chemical–genetic attenuation of focal neocortical seizures. *Nat. Commun.* 5, 1–9.
- Kujirai, K., Kujirai, T., Sinkjaer, T., Rothwell, J.C., 2006. Associative plasticity in human motor cortex during voluntary muscle contraction. *J. Neurophysiol.* 96, 1337–1346.
- Li, X., Yang, X., Sun, Z., 2020. Alpha rhythm slowing in a modified thalamo-cortico-thalamic model related with Alzheimer's disease. *PLoS One* 15, e0229950.
- Liang, Z., Duan, X., Su, C., Voss, L., Sleight, J., Li, X., 2015. A pharmacokinetics-neural mass model (PK-NMM) for the simulation of EEG activity during propofol anesthesia. *PLoS One* 10, e0145959.
- Liley, D., Walsh, M., 2013. The mesoscopic modeling of burst suppression during anesthesia. *Front. Comput. Neurosci.* 7, 46.
- Litvak, V., Jafarian, A., Zeidman, P., Tibon, R., Henson, R.N., Friston, K., 2019. There's no such thing as a 'true' model: the challenge of assessing face validity. In: 2019 IEEE International Conference on Systems, Man and Cybernetics (SMC). IEEE, pp. 4403–4408.
- Litvak, V., Mattout, J., Kiebel, S., Phillips, C., Henson, R., Kilner, J., Barnes, G., Oostenveld, R., Daunizeau, J., Flandin, G., 2011. EEG and MEG data analysis in SPM8. *Comput. Intell. Neurosci.* 2011.
- Liu, S., Ching, S., 2017. Homeostatic dynamics, hysteresis and synchronization in a low-dimensional model of burst suppression. *J. Math. Biol.* 74, 1011–1035.
- Lucarini, V., 2008. Response theory for equilibrium and non-equilibrium statistical mechanics: causality and generalized Kramers-Kronig relations. *J. Stat. Phys.* 131, 543–558.
- Lucarini, V., Speranza, A., Vitolo, R., 2007. Parametric smoothness and self-scaling of the statistical properties of a minimal climate model: what beyond the mean field theories? *Physica D* 234, 105–123.
- Marreiros, A.C., Kiebel, S.J., Daunizeau, J., Harrison, L.M., Friston, K.J., 2009. Population dynamics under the Laplace assumption. *Neuroimage* 44, 701–714.
- Mccarthy, M., Moore-Kochlacs, C., Gu, X., Boyden, E., Han, X., Kopell, N., 2011. Striatal origin of the pathologic beta oscillations in Parkinson's disease. *Proc. Natl. Acad. Sci.* 108, 11620–11625.
- Mina, F., Benquet, P., Pasnicu, A., Biraben, A., Wendling, F., 2013. Modulation of epileptic activity by deep brain stimulation: a model-based study of frequency-dependent effects. *Front. Comput. Neurosci.* 7, 94.
- Moody Jr, W.J., Futamachi, K.J., Prince, D.A., 1974. Extracellular potassium activity during epileptogenesis. *Exp. Neurol.* 42, 248–263.
- Moran, R., 2015. Deep brain stimulation for neurodegenerative disease: a computational blueprint using dynamic causal modeling. *Prog. Brain Res.* Elsevier.
- Moran, R.J., Kiebel, S.J., Stephan, K.E., Reilly, R., Daunizeau, J., Friston, K.J., 2007. A neural mass model of spectral responses in electrophysiology. *Neuroimage* 37, 706–720.
- Moran, R.J., Pinotsis, D.A., Friston, K.J., 2013. Neural masses and fields in dynamic causal modeling. *Front. Comput. Neurosci.* 7, 57.
- Moran, R.J., Stephan, K.E., Dolan, R.J., Friston, K.J., 2011. Consistent spectral predictors for dynamic causal models of steady-state responses. *Neuroimage* 55, 1694–1708.
- Muheim, C.M., Spinnler, A., Sartorius, T., Dürr, R., Huber, R., Kabagema, C., Ruth, P., Brown, S.A., 2019. Dynamic-and frequency-specific regulation of sleep oscillations by cortical potassium channels. *Curr. Biol.* 29, 2983–2992 e3.
- Nevado-Holgado, A.J., Marten, F., Richardson, M.P., Terry, J.R., 2012. Characterising the dynamics of EEG waveforms as the path through parameter space of a neural mass model: application to epilepsy seizure evolution. *Neuroimage* 59, 2374–2392.
- Panayiotopoulos, C.P., 2010. Atlas of epilepsies. Springer Science & Business Media.
- Papadopoulou, M., Cooray, G., Rosch, R., Moran, R., Marinazzo, D., Friston, K., 2017. Dynamic causal modelling of seizure activity in a rat model. *Neuroimage* 146, 518–532.
- Papadopoulou, M., Leite, M., Van Mierlo, P., Vonck, K., Lemieux, L., Friston, K., Marinazzo, D., 2015. Tracking slow modulations in synaptic gain using dynamic causal modelling: validation in epilepsy. *Neuroimage* 107, 117–126.
- Park, H.-J., Friston, K.J., Pae, C., Park, B., Razi, A., 2018. Dynamic effective connectivity in resting state fMRI. *Neuroimage* 180, 594–608.
- Park, H.-J., Pae, C., Friston, K., Jang, C., Razi, A., Zeidman, P., Chang, W.S., Chang, J.W., 2017. Hierarchical dynamic causal modeling of resting-state fMRI reveals longitudinal changes in effective connectivity in the motor system after thalamotomy for essential tremor. *Front. Neurol.* 8, 346.
- Penny, W.D., Friston, K.J., Ashburner, J.T., Kiebel, S.J., Nichols, T.E., 2011. Statistical Parametric Mapping: The Analysis of Functional Brain Images. Elsevier.
- Peron, S.P., Gabbiani, F., 2009. Role of spike-frequency adaptation in shaping neuronal response to dynamic stimuli. *Biol. Cybern.* 100, 505–520.
- Purdon, P.L., Pierce, E.T., Mukamel, E.A., Prerau, M.J., Walsh, J.L., Wong, K.F.K., Salazar-Gomez, A.F., Harrell, P.G., Sampson, A.L., Cimenser, A., 2013. Electroencephalogram signatures of loss and recovery of consciousness from propofol. *Proc. Natl. Acad. Sci.* 110, E1142–E1151.
- Rebola, N., Srikanth, B.N., Mülle, C., 2010. Activity-dependent synaptic plasticity of NMDA receptors. *J. Physiol.* 588, 93–99.
- Reick, C.H., 2002. Linear response of the Lorenz system. *Phys. Rev. E* 66, 036103.
- Rodriguez, R., Kallenbach, U., Singer, W., Munk, M.H., 2004. Short-and long-term effects of cholinergic modulation on gamma oscillations and response synchronization in the visual cortex. *J. Neurosci.* 24, 10369–10378.
- Rosa, M.J., Kilner, J.M., Penny, W.D., 2011. Bayesian comparison of neurovascular coupling models using EEG-fMRI. *PLoS Comput. Biol.* 7, e1002070.
- Rosch, R., Baldeweg, T., Moeller, F., Baier, G., 2018a. Network dynamics in the healthy and epileptic developing brain. *Netw. Neurosci.* 2, 41–59.
- Rosch, R.E., Hunter, P.R., Baldeweg, T., Friston, K.J., Meyer, M.P., 2018b. Calcium imaging and dynamic causal modelling reveal brain-wide changes in effective connectivity and synaptic dynamics during epileptic seizures. *PLoS Comput. Biol.* 14, e1006375.
- Rosch, R.E., Wright, S., Cooray, G., Papadopoulou, M., Goyal, S., Lim, M., Vincent, A., Upton, A.L., Baldeweg, T., Friston, K.J., 2018c. NMDA-receptor antibodies alter cortical microcircuit dynamics. *Proc. Natl. Acad. Sci.* 115, E9916–E9925.
- Rovainen, C., 1983. Generation of respiratory activity by the lamprey brain exposed to picrotoxin and strychnine, and weak synaptic inhibition in motoneurons. *Neuroscience* 10, 875–882.
- Ruelle, D., 2009. A review of linear response theory for general differentiable dynamical systems. *Nonlinearity* 22, 855.
- Saggio, M.L., Crisp, D., Scott, J.M., Karoly, P., Kuhlmann, L., Nakatani, M., Murai, T., Dümpelmann, M., Schulze-Bonhage, A., Ikeda, A., 2020. A taxonomy of seizure dynamotypes. *Elife* 9, e55632.
- Sanders, J.A., Verhulst, F., Murdock, J., 2007. Averaging Methods in Nonlinear Dynamical Systems. Springer.
- Schiff, S.J., 2012. Neural Control Engineering: The Emerging Intersection Between Control Theory And Neuroscience. MIT Press.
- Shaw, A., Moran, R.J., Muthukumaraswamy, S.D., Brealy, J., Linden, D., Friston, K.J., Singh, K.D., 2017. Neurophysiologically-informed markers of individual variability and pharmacological manipulation of human cortical gamma. *Neuroimage* 161, 19–31.
- Sherman, M.A., Lee, S., Law, R., Haegens, S., Thorn, C.A., Hämmäläinen, M.S., Moore, C.I., Jones, S.R., 2016. Neural mechanisms of transient neocortical beta rhythms: converging evidence from humans, computational modeling, monkeys, and mice. *Proc. Natl. Acad. Sci.* 113, E4885–E4894.
- Shin, H., Law, R., Tsutsui, S., Moore, C.I., Jones, S.R., 2017. The rate of transient beta frequency events predicts behavior across tasks and species. *Elife* 6, e29086.
- Spitzer, B., Haegens, S., 2017. Beyond the status quo: a role for beta oscillations in endogenous content (re) activation. *eneuro* 4.
- Steyn-Ross, A., Steyn-Ross, M., 2010. Modeling Phase Transitions in the Brain. Springer.
- Traub, R.D., Whittington, M.A., Buhl, E.H., Lebeau, F.E., Bibbig, A., Boyd, S., Cross, H., Baldeweg, T., 2001. A possible role for gap junctions in generation of very fast EEG oscillations preceding the onset of, and perhaps initiating, seizures. *Epilepsia* 42, 153–170.
- Van de steen, F., Almgren, H., Razi, A., Friston, K., Marinazzo, D., 2019. Dynamic causal modelling of fluctuating connectivity in resting-state EEG. *Neuroimage* 189, 476–484.
- Van Wijk, B.C., Cagnan, H., Litvak, V., Kühn, A.A., Friston, K.J., 2018. Generic dynamic causal modelling: an illustrative application to Parkinson's disease. *Neuroimage* 181, 818–830.
- Von Storch, H., Zwiers, F.W., 2001. Statistical Analysis in Climate Research. Cambridge university press.
- Voss, L.J., Brock, M., Carlsson, C., Steyn-Ross, A., Steyn-Ross, M., Sleight, J.W., 2012. Investigating paradoxical hysteresis effects in the mouse neocortical slice model. *Eur. J. Pharmacol.* 675, 26–31.
- Wei, Y., Ullah, G., Schiff, S.J., 2014. Unification of neuronal spikes, seizures, and spreading depression. *J. Neurosci.* 34, 11733–11743.
- Wendling, F., Bartolomei, F., Bellanger, J., Chauvel, P., 2002. Epileptic fast activity can be explained by a model of impaired GABAergic dendritic inhibition. *Eur. J. Neurosci.* 15, 1499–1508.
- Wendling, F., Benquet, P., Bartolomei, F., Jirsa, V., 2016. Computational models of epileptiform activity. *J. Neurosci. Methods* 260, 233–251.
- Wood, H., 2014. Chemical–genetic seizure silencing—unlocking the potential. *Nat. Rev. Neurol.* 10, 365–365.
- Zeidman, P., Jafarian, A., Corbin, N., Seghier, M.L., Razi, A., Price, C.J., Friston, K.J., 2019a. A guide to group effective connectivity analysis, part 1: first level analysis with DCM for fMRI. *Neuroimage* 200, 174–190.
- Zeidman, P., Jafarian, A., Seghier, M.L., Litvak, V., Cagnan, H., Price, C.J., Friston, K.J., 2019b. A guide to group effective connectivity analysis, part 2: Second level analysis with PEB. *Neuroimage* 200, 12–25.