



NETWORK NEUROSCIENCE

an open access  journal



Citation: Medrano, J., Friston, K., & Zeidman, P. (2024). Linking fast and slow: The case for generative models. *Network Neuroscience*. Advance publication. https://doi.org/10.1162/netn_a_00343

DOI:
https://doi.org/10.1162/netn_a_00343

Received: 8 August 2023
Accepted: 11 October 2023

Corresponding Author:
Johan Medrano
johan.medrano@ucl.ac.uk

Handling Editor:
Olaf Sporns

Copyright: © 2023
Massachusetts Institute of Technology
Published under a Creative Commons
Attribution 4.0 International
(CC BY 4.0) license



REVIEW

Linking fast and slow: The case for generative models

Johan Medrano , Karl Friston , and Peter Zeidman 

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

Keywords: Temporal scales, Generative models, Hierarchical modelling, Bayesian statistics, Dynamical systems, Hidden Markov models

ABSTRACT

A pervasive challenge in neuroscience is testing whether neuronal connectivity changes over time due to specific causes, such as stimuli, events, or clinical interventions. Recent hardware innovations and falling data storage costs enable longer, more naturalistic neuronal recordings. The implicit opportunity for understanding the self-organised brain calls for new analysis methods that link temporal scales: from the order of milliseconds over which neuronal dynamics evolve, to the order of minutes, days, or even years over which experimental observations unfold. This review article demonstrates how hierarchical generative models and Bayesian inference help to characterise neuronal activity across different time scales. Crucially, these methods go beyond describing statistical associations among observations and enable inference about underlying mechanisms. We offer an overview of fundamental concepts in state-space modeling and suggest a taxonomy for these methods. Additionally, we introduce key mathematical principles that underscore a separation of temporal scales, such as the slaving principle, and review Bayesian methods that are being used to test hypotheses about the brain with multiscale data. We hope that this review will serve as a useful primer for experimental and computational neuroscientists on the state of the art and current directions of travel in the complex systems modelling literature.

AUTHOR SUMMARY

Exploring changes in brain connectivity over time is a major challenge in neuroscience. This review article discusses the application of hierarchical generative models and Bayesian statistical methods to investigate modulators of neuronal dynamics across different temporal scales. By utilizing these innovative techniques, researchers can move beyond describing statistical associations and ask questions about underlying mechanisms. The article provides an overview of state-space modelling, dynamics, and a taxonomy of methods. It also introduces mathematical principles that link temporal scales and reviews the use of Bayesian statistical methods in testing hypotheses about the brain using multiscale data. This review aims to serve as a resource for experimental and computational neuroscientists by presenting the current state of the art and future directions of travel in modelling literature.

UNDERSTANDING THE BRAIN: A MULTISCALE PROBLEM

Phenomena that span multiple temporal scales are ubiquitous in brain research (Breakspear, 2017; D'Angelo & Jirsa, 2022). For instance, an objective of epilepsy research is to understand

the genesis of seizures by analysing the itinerancy of cortical circuitry parameters over seconds to minutes (Courtiol, Guye, Bartolomei, Petkoski, & Jirsa, 2020; Depannemaecker, Destexhe, Jirsa, & Bernard, 2021; Depannemaecker, Ezzati, Wang, Jirsa, & Bernard, 2023; V. Jirsa et al., 2023; V. K. Jirsa, Stacey, Quilichini, Ivanov, & Bernard, 2014; Lisi, Rivela, Takai, & Morimoto, 2018; Lopez-Sola et al., 2022; Ponce-Alvarez et al., 2015; R. E. Rosch, Hunter, Baldeweg, Friston, & Meyer, 2018). In naturalistic experiments, researchers study the dynamics of the freely-behaving brain in exchange with its environment, on the scale of minutes to hours (Demirtaş et al., 2019; Echeverria-Altuna et al., 2022; Lee, Aly, & Baldassano, 2021; Meer, Breakspear, Chang, Sonkusare, & Cocchi, 2020; Shain, Blank, van Schijndel, Schuler, & Fedorenko, 2020). On a larger time scale, studying the progression of Alzheimer's disease requires accounting for processes evolving over months, years, or even decades (Graham, Marr, Buss, Sullivan, & Fair, 2021; Wang et al., 2020). In these examples, a challenge is to understand how slow and fast components interact; in other words, how slowly changing variables shape the evolution of rapidly changing ones, and how the latter reciprocally influence the former (Ellis, Noble, & O'Connor, 2012).

This calls for a statistical framework for evaluating hypotheses when empirical observations span multiple temporal scales, or when the hypotheses are framed on a different time scale from that of the observations. Quantitatively comparing the evidence for hypotheses amounts to building generative models and comparing their ability to predict the data, in terms of model evidence (a.k.a., marginal likelihood). Modelling multiple temporal scales affords the potential to better predict future observations.

This review is based on the overarching idea that evaluating hypotheses with multiscale time series can be systematically addressed by constructing hierarchical models, where each level of the hierarchy accounts for a different temporal scale. Separating temporal scales can be rigorously justified and—by pairing such models with Bayesian procedures—they become useful tools for investigating the causes of observed data.

How can hierarchical models be used to capture changes in neuronal connectivity at different temporal scales, and what is the mathematical justification for their use? How to evaluate the evidence for competing hierarchical models, in order to address neuroscientific hypotheses or make modelling decisions (e.g., which temporal scales should be modelled)? These questions are the focus of this review. In the three sections that follow, we first introduce key concepts for modelling the dynamics of time series data at a single timescale. We then detail principles and strategies for separating timescales within mathematical models. Third, we introduce Bayesian statistical methods for evaluating the quality of multiscale models, before concluding with an empirical example in the field of epilepsy research.

MODELING TIME SERIES

To test hypotheses about the temporal structure of our observations, we need to model their generating process. Clearly, the form of the model depends on the nature of the data, for example, whether it is continuous or discrete, and on the assumed nature of the generating process, for example, deterministic or stochastic. This section introduces several key concepts in time series modelling, before considering a unified view of the different approaches that are suitable for neuroimaging. These concepts provide the foundation for linking temporal scales, detailed in the remainder of the article.

State-Space Models

When observations come from a continuous set of values, the modelling problem lends itself to the apparatus of dynamical systems. In particular, state-space models capture the evolution

over time of some latent or unobserved variables called states, \mathbf{x} , which give rise to observations \mathbf{y} . Using this framework, which was established in the 1960s in the field of control theory (Kalman, 1960), a broad class of dynamical systems can be modelled using equations of the form

$$\begin{cases} \frac{d}{dt}\mathbf{x}(t) &= f(\mathbf{x}(t), \mathbf{u}(t), \theta) \\ \mathbf{y}(t) &= g(\mathbf{x}(t), \mathbf{u}(t), \theta) + \omega(t) \end{cases} \quad (1)$$

where all variables can be vectors. The first equation is the evolution equation or state equation. At time t , the velocity of the system's states $\frac{d}{dt}\mathbf{x}(t)$ depends on the states themselves $\mathbf{x}(t)$, some inputs $\mathbf{u}(t)$, and some parameters θ , which can be viewed as states that change infinitely slowly. The states are observed through the observations $\mathbf{y}(t)$ produced by the observation equation. The observations might be corrupted by some additive measurement noise $\omega(t)$.

A state space is a geometric space, the axes of which are the states of the system (e.g., the firing rates of different neuronal populations), as illustrated in Figure 1. Given some inputs and parameters, the first line of Equation 1 attributes a velocity vector at each point of state space. The mapping from states to velocity is the flow of the system and plays an important role in understanding its evolution: every state trajectory evolves along the flow. The heterogeneity of the flow (intuitively, arrows pointing at each other) gives rise to attractors, that is, dense sets of points (i.e., a single point, a circle, or a sphere) attracting trajectories in their basin of attraction. Attractors define the steady state behaviour of the system: any trajectory in the basin of attraction evolves towards—and remains in—the attractor.

A multistable dynamical system has multiple attractors, thus its steady state depends on the attraction basin in which the system is initialized. Additionally, changes in inputs or parameters can cause bifurcations, that is, changes in the flow such that the topology of the attractor changes, for instance, from a fixed point to a limit cycle (Guckenheimer & Holmes, 1983). Bifurcations and multistability appear only in nonlinear dynamical systems and play a critical role in modelling biological systems (Breakspear, 2017; Deco & Jirsa, 2012; Haken, 1978; Haken, Kelso, & Bunz, 1985; V. Jirsa, 2020; Kelso, 2012; Ponce-Alvarez, Kringelbach, & Deco, 2023; Tognoli & Kelso, 2014).

When the evolution is purely deterministic as in Equation 1, one can represent the output of a dynamical system as a function of past inputs, without reference to internal states. In other words, the output signal can be derived from the inputs, without having to solve an **initial value problem**, that is, to integrate the equations through time. For linear systems, the dependence of outputs on past inputs is captured by a linear **response function** which, convolved with an input time series, gives the output time series. For nonlinear systems, a series of higher order response functions (a Volterra expansion) captures the dependence of outputs on past inputs (Fliess, Lamnabhi, & Lamnabhi-Lagarigue, 1983).

An example application of response functions in functional MRI (fMRI) research is the use of a Volterra formulation to summarise the dynamics of the blood oxygen level-dependent (BOLD) response (Friston, Mechelli, Turner, & Price, 2000). In this case, the first-order Volterra kernel corresponds to the change in the BOLD response due to a brief stimulus, and the second-order Volterra kernel quantifies how the BOLD response to one stimulus depends on the time elapsed since the previous stimulus (i.e., nonlinear effects of time). These Volterra kernels can be generated from the underlying state-space model. In practice, one is interested in finding the simplest model of the observations that can explain the most data (this is formally motivated thereafter). In fMRI, the relevant properties

Initial value problem:
The problem of solving a differential equation with given initial conditions.

Response function:
A function that describes the output or behaviour of a system in response to a given input or stimulus.

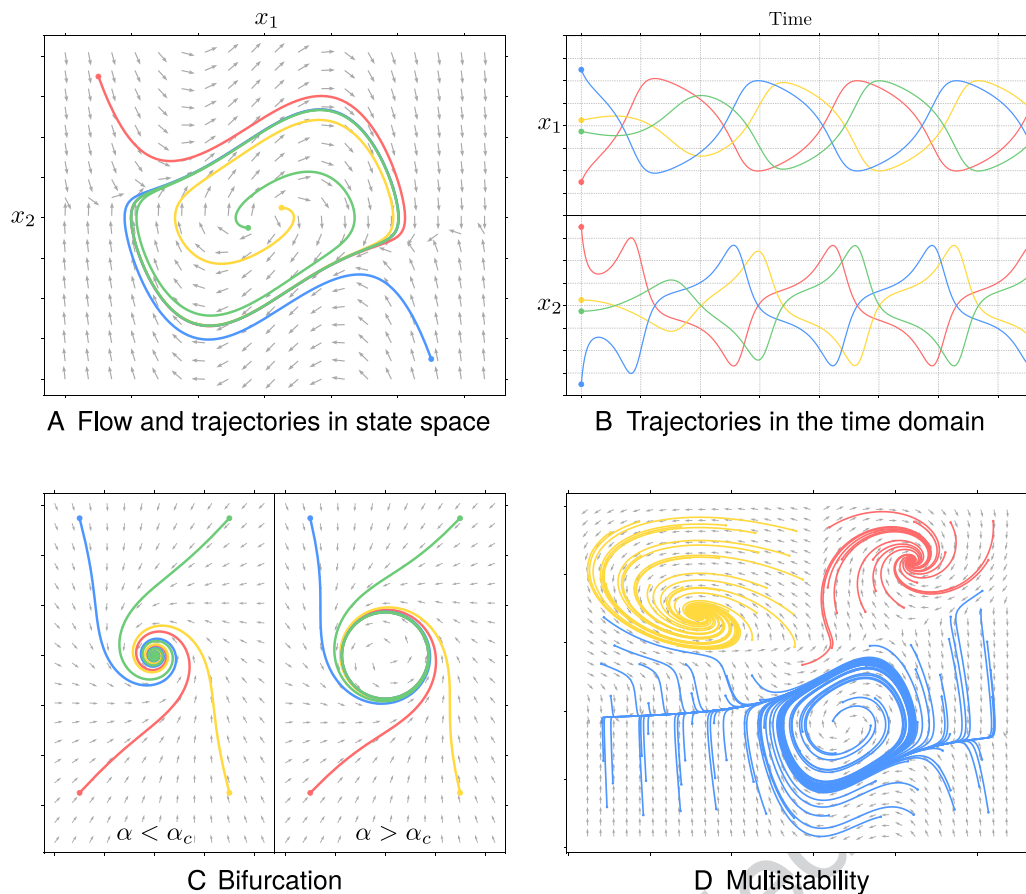


Figure 1. Important concepts of the theory of dynamical systems. (A) An example of flow in state-space (grey arrows), governing the evolution of trajectories (coloured curves) from different initial states (coloured circles). (B) The corresponding trajectories in the time domain for both x_1 and x_2 axes. (C) An example of bifurcation of the flow: trajectories converge towards a fixed point of state space when the bifurcation parameter α is below a critical value α_c , and towards a limit cycle when the bifurcation parameter is above the critical value (Andronov-Hopf bifurcation). (D) An example of a multistable system: the attractor to which the trajectory evolves depends on the initial state, as indicated by the colours of the trajectories.

of nonlinear systems are well-captured by retaining the first terms of the Volterra expansion, which corresponds to approximating the flow of the nonlinear dynamical system with a bilinear (Friston, Harrison, & Penny, 2003; Friston et al., 2000) or second-order polynomial form (Stephan et al., 2008).

For modelling neuronal connectivity and dynamics, the response function of a dynamical system plays an important role, because it captures how the system transforms the frequency spectrum of the input signal. This is particularly helpful for modelling systems that are driven by random fluctuations, such as neural mass models, which capture the activity of populations of neurons. When driven by endogenous cortical noise, it is impossible to track the evolution of a neural mass model: despite having a deterministic evolution, their trajectories depend nonlinearly on a particular realisation of the input noise process which is unavailable to the experimenter. In that case, we can model the output spectrum of neural masses by applying their linearized response function to their input spectrum (Chen, Kiebel, & Friston, 2008; R. Moran, Pinotsis, & Friston, 2013; R. J. Moran et al., 2007; R. J. Moran, Stephan, Dolan, & Friston, 2011; R. E. Rosch, Wright, et al.,

Effective connectivity:
The directed influence of neuronal populations on one another.

2018; Symmonds et al., 2018). The use of a linear approximation can be motivated by the fact that coupled neural systems exhibit linear dynamics over short timescales (Heitmann & Breakspear, 2018).

Dynamic causal modelling (DCM) (Frässle, Aponte, et al., 2021; Friston et al., 2003) and The Virtual Brain (V. K. Jirsa et al., 2017; Sanz Leon et al., 2013) are examples of applying state-space modelling to neuronal dynamics, and in particular for investigating **effective connectivity**, namely, the effects of neuronal populations on one another. In this context, the states are continuous variables, and the evolution equation typically describes the change in neuronal activity over time, whereas the observation function generates the measurements one would expect to make, for example, using electro-encephalography (EEG) or functional magnetic resonance imaging (fMRI).

A Taxonomy of Modelling Frameworks

Forming a model necessitates two key decisions. First, whether the dynamics will be deterministic, as in Equation 1, or stochastic, meaning there will be some degree of randomness to the evolution of the states. Stochastic models are covered by the more general framework of random dynamical systems (Arnold, 1995). The second decision is whether the states themselves will be discrete or continuous variables. Examples of these different kinds of modelling frameworks are provided in Figure 2.

A special case arises when the dynamics are stochastic, and the states are discrete values. This typically calls for Markov chains and related models (Fraser, 2008). Markov chains are stochastic processes that model probabilistic transitions between a finite collection of states. More precisely, Markov chains model the evolution of the states' distribution, that is, the probability of being in a certain state at a certain time, by means of state transition probabilities,

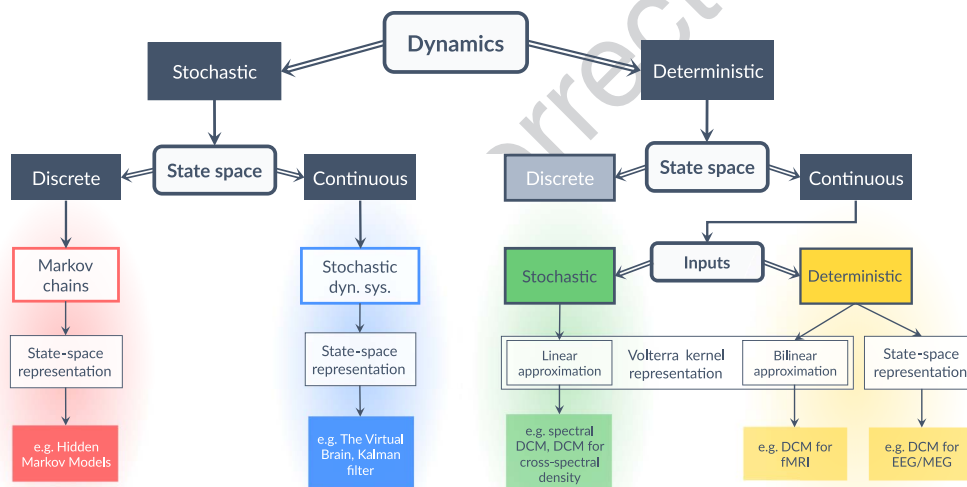


Figure 2. Taxonomy of the different modelling frameworks discussed here. The key factors guiding the selection of a particular framework are the nature of dynamics (discrete or continuous) and the nature of state space (stochastic or deterministic). In addition, the nature of the inputs (stochastic or deterministic) has relevance for continuous deterministic systems. In the particular case of continuous deterministic systems with stochastic inputs, one can use a linear response function, which is the first-order term of the Volterra kernel representation of the system, to directly approximate the outputs from the inputs without reference to the states. Effectively, this implies that the dynamics do not need to be integrated over time, which greatly simplifies model inversion. In all other cases, model inversion entails tracking the states or their distribution through time.

Hidden Markov model:
A probabilistic model that describes a sequence of observable events generated by an underlying unobservable Markov chain.

which gives the probability of switching from a state to another. Markov chains fulfil the Markov property: knowing the current state distribution and the state transition probabilities is sufficient to determine the future state distribution.

Markov chains have the same role as the evolution equation of dynamical systems. As such, they can also be equipped with an observation function transforming abstract “states” into the expected distribution over observations. The resulting models are known as **Hidden Markov Models** (HMMs) and are popular devices in modelling time series (Baker et al., 2014; Bishop & Nasrabadi, 2006; Cabral et al., 2017; Rabiner, 1989; Vidaurre et al., 2016; Vidaurre, Smith, & Woolrich, 2017). Relevant examples of HMM applications include (i) analysing the switching dynamics of resting-state networks in a large cohort (Vidaurre et al., 2018), (ii) evaluating how brain network dynamics are altered during natural movie watching (Meer et al., 2020), and (iii) identifying brain networks activated during replay (Higgins et al., 2021).

Both stochastic and deterministic state-space models are routinely used in the DCM modelling framework, which is implemented in freely available software (SPM; Penny, Friston, Ashburner, Kiebel, & Nichols, 2011; TAPAS, Frässle, Aponte, et al., 2021). DCM pairs the specification of models with Bayesian system identification techniques, in order to estimate the evidence for alternative candidate models. Recent developments have included continuous state-space models of psychiatric disease progression (Friston, Redish, & Gordon, 2017), whole-brain effective connectivity in resting-state fMRI (Frässle, Harrison, et al., 2021; Frässle et al., 2018), as well as discrete state-space models of Covid-19 progression (Friston, Flandin, & Razi, 2022; Friston et al., 2020). Figure 2 includes examples of particular variants of DCM that handle different kinds of states and evolution equations.

To summarize, time series can be modelled as resulting from the dynamics of latent states, where the dynamics may be deterministic or stochastic and where the states may be continuous or discrete. Despite the apparent heterogeneity of modelling methods, they are constructed from common components: an evolution equation, guiding the system’s state through time, and an observation equation, producing the measured quantities. Response functions can then be derived from the state-space model, serving as a useful device for obtaining the output of a system from its input. The choice of approach is determined only by the nature of the states and that of the evolution equation, for the particular application domain.

SEPARATING TIMESCALES

The fundamental concepts in time series modelling reviewed above can be applied to situations where observations have a temporal structure at different temporal scales. This is the case when the duration of the observations is long enough to allow the slower temporal features to unfold—and when the sampling rate is high enough to capture fast temporal features. In this case, modelling the evolution at different time scales becomes crucial, in order to infer the common underlying generative processes or mechanisms that give rise to the data.

Motivation

Considering the temporal evolution of time series at various scales is a natural approach for studying neuronal dynamics. One common practice involves examining power fluctuations in the frequency bands of EEG or MEG data. The power spectrum of a signal summarizes short-term signal changes, while the power of specific frequency bands (a.k.a., band power) reveals longer term variations. By using a model with parameters governing the shape of the spectral density, such as the response function of a neural mass model, we can connect the rapid

fluctuations in electrical activity over milliseconds to the slow evolution of these parameters over seconds or minutes.

These models allow us to quantify how experimental manipulations relate to continuous changes in slow parameters. For instance, R. E. Rosch, Wright, et al. (2018) demonstrated how the concentration of an epilepsy-inducing drug modulated the power spectra of neuronal activity by affecting certain synaptic rate constants. In the future, this type of model could prove valuable in characterising the dynamics of the self-inhibition of superficial pyramidal neurons, thought to encode prediction precision in predictive coding accounts of brain function, for instance, during naturalistic experiences (Adams, Bauer, Pinotsis, & Friston, 2016; Bastos et al., 2012; Kanai, Komura, Shipp, & Friston, 2015).

However, building and analysing models with dynamics over different temporal scales can be challenging. Determining the appropriate temporal scale for each variable—and understanding its interaction with other variables—is not straightforward. In many cases, interactions among variables evolving at different timescales can result in circular causality, complicating the analysis and leading to complex multiscale dynamics. A key example of this is experience-dependent plasticity, where slow fluctuations in synaptic efficacy depend upon fast fluctuations in neuronal activity. At the same time, distributed neuronal activity depends upon synaptic efficacy. Fortunately, there are mathematical principles that can help simplify these complex models while still capturing the essential aspects of their dynamics.

This section proceeds in three parts. First, we introduce mathematical arguments, starting with foundational and abstract concepts and gradually transition to more practical applications. Second, we explore the application of these arguments to stochastic dynamics. Last, we discuss their relevance to models with slow discrete dynamics such as HMMs.

Mathematical Perspectives

In dynamical systems, different temporal scales appear when some states evolve quickly relative to others. That is, if each state is conceptualized as forming a dimension of a state space (as we saw in Figure 1A), then different temporal scales emerge when the flow governing the evolution of trajectories is stronger in some directions than in others. The components of the dynamics in the stronger directions of the flow converge rapidly towards their steady state, with negligible displacements along the weaker directions of the flow. From the perspective of fast components, slow components may be considered as, effectively, static. Reciprocally, on the temporal scale at which slow components evolve, the fast components instantaneously reach their steady state. Inherent to the concept of separation of temporal scales is the reduction of the number of dynamical degrees of freedoms: over long periods of time, the macroscopic behaviour of the system can be described by the evolution of slow variables (Kuramoto & Nakao, 2019). This natural separation of temporal scales is crucial for modelling high-dimensional systems such as the brain (see Figure 3).

Mathematically, the separation of temporal scales in a nonlinear dynamical system can be proved rigorously. The centre manifold theorem (see Figure 4) is used to show the exponential decay of the fast components towards a manifold, that is, a generalised surface that is aligned with the weak directions of the flow (Carr, 1981). When the centre manifold theorem applies, one can consider that the dynamics are separated into fast and slow components evolving on different temporal scales, and approximate the overall dynamics by the trajectory on the manifold. However, this approximation is a local result; in other words, it can only be used in the vicinity of a fixed point. Recently, interest has been drawn to inertial manifolds. Similarly to centre manifolds, inertial manifolds are surfaces that attract trajectories exponentially quickly,

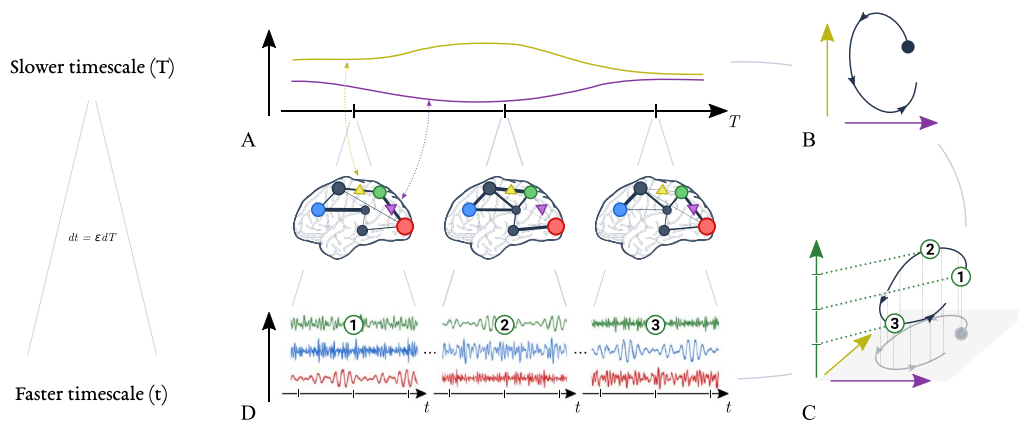


Figure 3. Multiscale dynamics of brain signals: mapping slow and fast variables. (A) Slow quantities in the brain, such as synaptic efficacy between regions, exhibit large time constants and evolve slowly over time. The evolution of two slow variables are illustrated here, as yellow and purple lines. (B) The evolution of slow variables can also be represented by dynamics in a slow state space. (C) Importantly, for every location in the slow state space (horizontal axes), there is a corresponding mode of fast dynamics (vertical axis). Three modes are depicted here, numbered 1–3. These fast dynamics give rise to rapid brain signals, such as field potentials in pyramidal neurons (D). The mathematical relationship between the slow and the fast timescale is given by $dt = \epsilon dT$ ($\epsilon \ll 1$); in other words, the dynamics at faster scale t unfolds over a fraction (ϵ) of the slower scale T . In summary, the brain is understood to navigate slowly (A) through a repertoire of fast stable dynamics (D). Crucially, the slow variables are directly linked to the dynamics of the fast variables (C). Similarly, changes in the fast variables’ dynamics can be attributed to changes in the slow variables. Therefore, modelling the complex dynamics of multiscale dynamical systems can be simplified by focusing on the dynamics of the slow variables and the mapping from slow to fast variables.

and can indeed be seen as *global* centre manifolds (Foiás, Sell, & Temem, 1988). However, in contrast to centre manifolds, there is no generic method to prove the existence of an inertial manifold, which needs to be assessed on a case by case basis.

Manifolds give a practical way to construct multiscale dynamical systems: one can specify an equation for the (inertial) manifold, mapping from slow to fast states, and prescribe strong normal (perpendicular) flow to make fast states quickly collapse onto the manifold. The framework of structured flows on manifolds (SFM) adopts this procedure. SFM explains how slow low-dimensional variables, such as movement parameters, could be instantiated by the collective behaviour of a large number of fast variables, such as networks of neurons (V. Jirsa & Sheheitli, 2022; V. K. Jirsa, McIntosh, & Huys, 2019; Pillai & Jirsa, 2017). In SFM, the fast components of the flow drive neuronal activity to rapidly evolve into synchronization modes (functional modes) that correspond to points on the manifold. Simultaneously, the flow on the manifold progressively modifies the functional modes, thereby generating slowly changing behavioural variables.

Frameworks such as SFM are important for brain research because they give a mechanistic account of brain activity and its relationship to behaviour (V. Jirsa & Sheheitli, 2022; McIntosh & Jirsa, 2019). Interestingly, this type of theoretical modelling approach aligns well with the needs of empirical work on intracranial recordings of neuronal activity. Typically, an objective is to understand how movement variables are encoded by stable low-dimensional dynamics of neuronal populations in the motor cortex (Chaudhuri, Gerçek, Pandey, Peyrache, & Fiete, 2019; Churchland et al., 2012; Gallego et al., 2018; Langdon, Genkin, & Engel, 2023; Lara, Cunningham, & Churchland, 2018; Saxena, Russo, Cunningham, & Churchland, 2022). Combining more advanced modelling frameworks—together with the statistical machinery described later—can provide valuable tools to show on how empirical recordings support different theoretical perspectives on motor control.

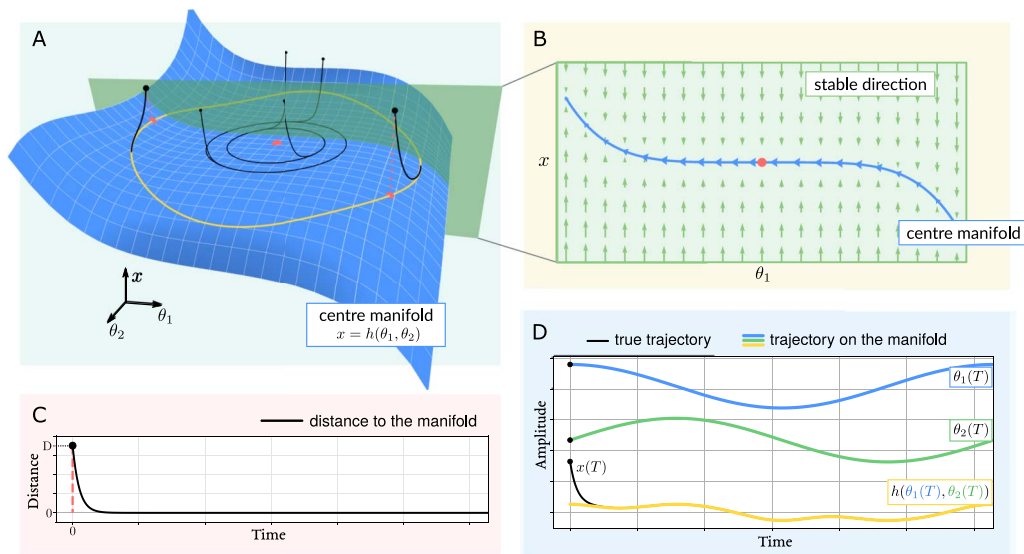


Figure 4. Illustration of the centre manifold theorem with a three-dimensional dynamical system. (A) The three-dimensional state-space of the system. The blue surface is the centre manifold, and gives a height $x = h(\theta_1, \theta_2)$ to each point of the (θ_1, θ_2) plane. Trajectories initialized away from the manifold converge rapidly towards the surface (black curves). This is due to the presence of a strong flow orthogonal to the centre manifold (B) for a section of state space. The strong flow (green arrows) converges towards the centre manifold (blue curve). The flow parallel to the manifold (blue arrows) is weaker by orders of magnitude. Hence, trajectories quickly collapse to the centre manifold before evolving alongside it. This is reflected in the exponential decay of the distance to the manifold (C). Hence, the x -component of the trajectory is well approximated by a static function from the location on the (θ_1, θ_2) plane (D). This can motivate an adiabatic approximation: as the rapidly changing x component of the trajectory can be approximated by the mapping $h(\theta_1, \theta_2)$, we may consider x as a spurious dimension of the system and restrict our description to the evolution on the (θ_1, θ_2) plane; in other words, we can approximate the fast vanishing states by a fixed mapping from the slow states.

Adiabatic approximation:
Neglecting rapidly changing variables in comparison to slowly changing ones.

Synergetics:
The study of complex systems and their emergent properties, examining interactions between parts to understand their holistic dynamics.

We can also go one step further with the separation of temporal scales and neglect the dynamics of fast variables when considering the evolution of slow variables (as illustrated in Figure 4). This separation is formally known as the **adiabatic approximation**: the fast variables are replaced by their steady-state solution, leaving only the differential equations of the slow variables (Friston, Li, Daunizeau, & Stephan, 2011; Haken, 1978). The adiabatic approximation is particularly appealing for multiscale systems as it allows one to replace references to the fast variables in the slow dynamics by a fixed mapping from slow variables, thereby finessing the problem of circular causality.

The mathematical arguments of slow-fast decomposition originate from theoretical physics and have been used in works on self-organisation and **synergetics** (Haken, 1977, 1978). Haken and colleagues framed the decomposition of temporal scales in terms of a “slaving principle,” where slow collective parameters, the order parameters, govern the dynamics of fast individual variables. This was elegantly used to construct the Haken-Kelso-Bunz (HKB) model of synchronisation during rhythmic finger movements, which provides a simple yet representative example of the slaving principle underpinning the slow-fast decomposition (Haken et al., 1985). They found that when participants moved their index fingers at a fast speed, those initially moving in opposite phases suddenly synchronized. By using the slaving principle, they summarised the complex behaviour of the coupled finger dynamics through a single slow collective variable (the phase difference), which served as an order parameter. Their model shows that when the speed of finger movements exceeds a critical frequency, a bifurcation occurs in the order parameter: the antiphase movements become unstable, and the fingers

synchronize. This celebrated example illustrates the power of the slaving principle in simplifying intricate multiscale phenomena.

To summarize, a dynamical system evolving on two different timescales can be described by a weak flow on a manifold, that is, a surface within state space, and a component normal to the manifold whose amplitude decays exponentially with time. This fact can be used, not only to assess the existence of a separation of temporal scales in a given system, but also to construct a multiscale dynamical system by prescribing a manifold, a flow on the manifold, and a strong convergent flow normal to the manifold. This is the approach taken by the SFM framework. Multiscale dynamical systems can be further simplified by discarding the normal component, a step known as the adiabatic approximation. The adiabatic approximation has a particular advantage when dealing with dynamical systems driven by noisy inputs.

Stochastic Dynamics and Conditioning

We have discussed separating timescales in continuous deterministic dynamical systems. However, the dynamics may be influenced by noise, and it is useful to understand how random fluctuations enter a dynamical system with multiple timescales.

A key mathematical result is that the centre manifold theory and the adiabatic approximation apply to stochastic systems (Boxler, 1989; Knobloch & Wiesenfeld, 1983). Naturally, the timescale separation takes a different form than in the deterministic case: we need to describe the stochastic evolution of the system over time, for instance, through the evolution of the state *probability distribution* (see Roberts, 2014, Chap. 21, for details). Under a separation of temporal scales, the state distribution factorizes into a marginal, time-dependent distribution of slow modes $p(r,t)$ and a conditional, time-independent distribution of fast components $p(s|r)$. In other words, the distribution of the fast modes is fully determined by the slow modes.

The application of the centre manifold theorem to stochastic systems has important consequences. The distribution of fast modes is conditioned on the slow modes. For multiscale stochastic dynamical systems, this conditioning separates the different temporal scales. Thus, a generative model of multiscale time series naturally has a hierarchical structure, where the hierarchical depth accommodates the multiple fine or coarse graining of the systems temporal structure.

Discrete Systems and the Slaving Principle

The slaving principle describes the general idea that slow quantities determine the evolution of fast quantities. For the particular case of continuous dynamical systems, the principle is evinced by the centre manifold theorem and the adiabatic approximation. However, the slaving principle is more general and also integrates common applications of discrete switching models such as HMMs. Here, we show how discrete systems can be reconciled with the overarching slaving principle.

When studying brain signals, it appears that brain activity is confined to modes or states of activity. These modes are defined by specific patterns of activity across the brain, which can be described through functional or effective connectivity between different brain regions. Notably, these stable modes of activity are activated in a sequential manner, and analysing the sequences of activation can provide valuable insights into the overall dynamics of the brain on a large scale. Analysing brain network dynamics—especially in resting state—has proved useful, for instance, in pain research, in epilepsy research (V. K. Jirsa et al., 2017; R. Rosch,

Baldeweg, Moeller, & Baier, 2018), and in Alzheimer's research (Kuang et al., 2019; Núñez et al., 2021; Smailovic et al., 2019).

A recent trend in neuroimaging involves employing HMMs to capture the sequential activation of the fast stable modes (Baker et al., 2014; Cabral et al., 2017; Vidaurre et al., 2016, 2017). With electrophysiological data, the overall evolution of brain-wide neuronal activity is represented by both the cross-spectral coherence, which captures continuous short-term brain activity, and a Markov chain that models the switching dynamics over longer time scales (O'Neill et al., 2018; Vidaurre et al., 2018, 2016). The use of HMMs is licensed by the fact that the transitions between states happen rapidly, whereas the duration spent in each state is sufficiently long to enable the rapid activity to reach its steady state (Deco & Jirsa, 2012; Ponce-Alvarez et al., 2015; Rabinovich, Huerta, Varona, & Afraimovich, 2008). Thus, the switching dynamics of neuronal connectivity in HMMs adheres to the slaving principle: the slow Markov chain governs the dynamics of the fast brain signals, as captured by the cross-spectral coherence.

To summarize, time series at different temporal scales can be separated and individually modelled. Separating temporal scales implies committing to the slaving principle, that is, the idea that slow quantities govern the dynamics of fast quantities. The mathematical arguments underpinning the slaving principle are, for now, restricted to continuous systems; however, the principle can be intuitively applied to systems with switching dynamics such as HMMs. Naturally, this perspective shows that a model of multiscale time series has a hierarchical structure, with each layer accounting for a different temporal scale. As a result, we can model the power spectra of EEG or MEG signals using a fast neural mass model alongside a slow model for parameter dynamics. Similarly, the dynamics of a group of neurons can be captured by a fast spike rate model along with a slow model for population dynamics. The next section will focus on how to utilize hierarchical models to test hypotheses with empirical data.

EVALUATING HYPOTHESES WITH HIERARCHICAL MODELS OF MULTISCALE TIME SERIES

A mathematical model, of the sort described above, specifies a hypothesis about how the measured data were generated. Arbitrating between hypotheses entails specifying a suitable set of models, and having the statistical machinery in place to compare their evidence. This section sets out recent developments in Bayesian statistical methods, which are proving useful in neuroimaging and related fields that deal with multiscale time series data and dynamic models.

Statistical Methodology and the Case for Modelling

In Bayesian statistics, the key quantity summarising the quality of a model is the log model evidence, also called the marginal likelihood, $\ln p(y|m)$. This is the log of the probability of having observed the data y given the model m . Two or more models can then be compared in terms of their relative log evidence, which is called the log **Bayes factor**: $\ln B = \ln p(y|m_1) - \ln p(y|m_2)$ (Kass & Raftery, 1995). This procedure is called Bayesian model comparison.

During model development, Bayesian model comparison provides a principled framework for evaluating modelling decisions. Then, when testing hypotheses in an empirical setting, Bayesian model comparison serves as the basis for evaluating how well each model explains the data. This widespread use of the log evidence is justified because it has a number of attractive properties. For instance, it can be decomposed into the difference between the accuracy

Bayes factor:

A statistical measure that quantifies the evidence supporting one hypothesis, as represented by a model, over another.

Posterior distribution:
The updated probability distribution
of a variable after considering new
data.

and complexity of a model, so it can be used to identify the simplest explanation for the data that maximizes the explained variance. Furthermore, through a generalization of the Neyman-Pearson lemma (Fowlie, 2021; Neyman & Pearson, 1933), it can be shown that the Bayes factor maximizes statistical power.

Generative models that describe how neuronal dynamics give rise to observed neuroimaging data typically contain many parameters. These may include coupling strengths between brain regions or the rate of depolarization of a neuronal membrane. Prior knowledge about model parameters are generally available, and the probability distribution over the parameters is refined by observing the data. Bayesian inversion refers to the process of transforming a prior probability distribution over the parameters into a **posterior distribution** through observing the data. This depends on calculating or estimating the model evidence, which as stated above, quantifies how well the model explains the data, once any uncertainty about the parameters has been accounted for.

Calculating the exact posterior distribution is generally impossible. Bayes rule gives the posterior distribution of parameters $p(\theta|y)$ as the joint probability of parameters and data $p(\theta, y)$ over the marginal probability of the data (or model evidence) $p(y)$, that is, $p(\theta|y) = p(\theta, y)/p(y)$. Computing the evidence requires marginalising out (i.e., integrating over) the parameters in the joint distribution (i.e., $p(y) = \int p(\theta, y)d\theta$), which is unfeasible in practice. Hence, one resorts to approximating the posterior, for which two categories of methods exist. With the first category of methods—Markov chain Monte Carlo—one first constructs a Markov chain whose equilibrium distribution is the posterior distribution and then samples from the Markov chain to approximate the posterior. With the second category of methods—variational Bayes—one first specifies the functional form of an approximate posterior and then minimises a statistical difference between the true and approximate posterior.

Variational methods are preferred when the objective is to evaluate hypotheses by comparing models. Although sampling methods can be used to compute any kind of complex, multimodal posterior probability distribution, there is no standardized method to reliably compute the model evidence. In contrast, variational Bayes methods approximate the log model evidence using a quantity known as the variational free energy or an evidence lower bound (ELBO). An optimization algorithm is applied, which searches for a setting of the model parameters that maximize the free energy. When maximized, the free energy approaches the log evidence. The quality of the approximation depends on the model complexity and becomes exact for general linear models. The variational Bayes treatment of stochastic and deterministic nonlinear dynamical systems has a long history of use in neuroimaging (Daunizeau, Friston, & Kiebel, 2009; Friston et al., 2003).

Inverting Hierarchical Models

We have established that evaluating hypotheses amounts to comparing models in terms of the ratio of model evidence, and that models of multiscale time series have a hierarchical structure, with higher levels of the hierarchy accommodating longer temporal scales. This raises the following question: how to invert hierarchical models and estimate their model evidence?

To recap, each level of a hierarchical model has parameters that are constrained by the level above. For example, with neuronal recordings, a neural mass model may be used at the first level, with parameters that encode synaptic connection strengths. Slow changes in these synaptic parameters over a longer time period (e.g., seconds or minutes) are described by the second level of the model, with parameters that govern the slow dynamics.

Empirical priors:

Prior distributions estimated from the empirical distribution of data or lower-level posteriors.

Bayesian model reduction:

A statistical technique based on Bayesian inference that simplifies complex models by eliminating irrelevant variables and retaining important ones.

The Bayesian treatment of hierarchical models is particular. Because it is generally hard to prescribe a priori the probability distribution of the parameters at arbitrary levels of the hierarchy, one is required to use **empirical priors**. The posterior distributions at each level are constrained by the priors of the level above. However, using an empirical Bayes approach with hierarchical models creates a circular dependency: empirical priors must be estimated from the posterior distribution of levels below, which in turn depends on empirical priors of levels above. An established solution to the circularity problem is to use an iterative procedure and alternately update our estimates of empirical priors and posteriors (Efron & Morris, 1973; Kass & Steffey, 1989).

A good example of empirical Bayes is found in Bayesian approaches to group studies. Subject-level data are modelled with parameters that are constrained by a group-level prior distribution. Each subject's model is inverted using the group level prior to produce a posterior estimate per subject. The prior group-level distribution is then refined from the collection of subject-level posterior estimates. The refined group-level empirical prior is used to produce new subject-level posterior estimates, which can be used to refine the empirical prior, and so forth.

Bayesian model reduction is a simpler approach that avoids the need to iterate between levels of the hierarchical model and is used extensively in the context of neuroimaging (Friston et al., 2016). The free energy—applied here to approximate the log evidence of the entire hierarchical model—can be separated into a partial free-energy for each hierarchical level. Each partial free-energy quantifies the complexity of that level and its accuracy at predicting the posterior density of the level below. Crucially, the free-energy of a level can be optimized regardless of higher levels. Thus, one can approximate the posterior of the first level using the empirical data, then approximate the posterior of the second level using the posterior of the first level, and repeat this procedure upwards in the hierarchy. This provides a straightforward analysis procedure for group studies in neuroimaging, which begins with modelling each participant's data individually (first level analysis), and then conveying the free energy and posteriors to the group level (second level analysis).

The Bayesian model reduction approach to hierarchical models echoes the separation of temporal scales (Jafarian, Zeidman, Wykes, Walker, & Friston, 2021). In hierarchical models of multiscale time series, each level represents a different temporal scale, with higher levels representing longer timescales. With a genuine empirical Bayes procedure, one would iterate over different timescales, which would make the problem computationally demanding. However, when we commit to the variational Bayes approach, each temporal scale is inverted only once.

In summary, the variational Bayes approach to inverting hierarchical temporal models is a statistical framework that can be used to quantitatively evaluate hypotheses about multiscale data. In particular, it extends a summary statistic approach, in which one would compute summary statistics of the fast variables and then perform post hoc analysis of the slow dynamics. In the following section, we present an example of how this treatment of multiscale dynamical systems has been used successfully.

EXAMPLE: ADIABATIC DYNAMIC CAUSAL MODELLING

We have now introduced three key ingredients for modelling time series data at multiple temporal scales: (i) state-space models, (ii) their extension to multiple temporal scales via hierarchical models, and (iii) Bayesian inference methods needed to invert these models and evaluate their evidence.

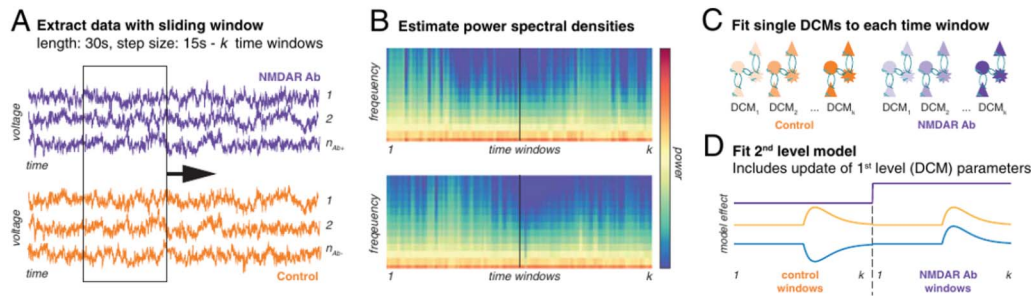


Figure 5. Hierarchical modelling approach used to link slow effects to fast observations. First, the authors extracted sliding window data from their LFP time series (A). Then, they estimated the power spectral density for each window, to produce a time-frequency representation of the data (B). Then they fitted a state-space model, called a canonical microcircuit (CMC) dynamic causal model (DCM), for each window of power spectral densities (C). This resulted in a time course of posterior densities for the parameters of the DCM models. Finally, the authors added a second level to the model to test for between-window effects, enabling them to evaluate hypotheses of interest, that is, the interaction between interventions (PTZ concentration, presence or absence of NMDAR-Ab) and the parameters of the CMC model (D). Adapted with permission from R. E. Rosch et al. (2018).

These methods are already proving useful for modelling multiscale data in practice. For instance, to investigate the role of antibodies against NMDA receptors (NMDAr-Ab) in autoimmune encephalitis, Rosch and colleagues used multiscale modelling to understand how NMDAr-Ab cause abnormal EEG spectra (R. E. Rosch, Wright, et al., 2018) (see Figure 5). They recorded the local field potential (LFP) of control and NMDAr-Ab-positive mouse models during epileptic seizures induced by pentylentetrazol (PTZ). The spectra of the LFPs were modelled for short time windows using a canonical microcircuit (CMC) model, composed of several interacting neuronal populations.

In their work, Rosch and colleagues used the CMC model as the first-level model (encapsulating fast variables) for each time window. The time constants of each neuronal population were the slow variables. The study evaluated how these slow variables interacted with the PTZ concentration, changing slowly through time, and with the presence of NMDAr-Ab. This was done by comparing models at the second level (i.e., between time windows). Under their model, they observed a strong effect of PTZ on the synaptic time constants of different neuronal populations that was further amplified by the presence of NMDAr-Ab.

This example illustrates how Bayesian model reduction and hierarchical temporal models can be used in practice. The complex multiscale modelling problem was decomposed in two simpler modelling problems:

- i. linking slow parameters of the models to the fast measurements,
- ii. modelling the effect of slow experimental effects on the slow parameters.

Opting for this decomposition allowed the authors to evaluate their research hypotheses in a principled manner. The method used has been more extensively formalized in the adiabatic DCM framework (Jafarian et al., 2021). The reader interested in an analogous method with discrete slow dynamics can refer to the dynamic effective connectivity framework introduced in Zarghami and Friston (2020).

CONCLUDING REMARKS AND FUTURE PERSPECTIVES

The mechanisms of interest for brain research generally span multiple temporal scales. Because of this apparent complexity, here we returned to first principles and pursued

established mathematical paths. Evaluating hypotheses amounts to constructing generative models of our observations, for which there are readily available methods. The slaving principle means that slow and fast processes can be analysed separately: we can therefore model each timescale individually, lending a hierarchical structure to our generative models. Moreover, the slaving principle entails switching dynamics, for instance, of brain networks dynamics, which licenses the combination of discrete models (e.g., HMMs) for slow changes in brain states and continuous models of fast neuronal dynamics, as we have illustrated.

Evaluating and comparing the evidence for hypotheses regarding multiscale time series requires us to invert hierarchical temporal models, that is, to compute the posterior distribution over model parameters. In most cases, the inversion problem cannot be solved exactly because of both the complexity and hierarchical structure of the model. This forces us to approximate the posterior distribution, a task for which variational Bayes methods are preferred. In particular, variational Bayes can be combined with empirical Bayes to invert arbitrary hierarchical models with empirical priors at intermediate levels. This suggests a simple procedure to invert hierarchical temporal models: first, model the fast temporal scale by fitting a model to each time window of time series data, and then invert these models to obtain a time series of estimated parameters (posteriors), which are expected to change slowly. Finally, use the time series of these slowly changing parameters as observations for the level above in the hierarchy.

A practical question is how to determine whether it is better to model several time scales, and how to identify which time scales shall be modelled. The general answer is that additional model complexity must be justified by having a better explanation of the data. The trade-off between model complexity and accuracy is effectively summarized by the log evidence of the model—and its approximation, the variational free energy—(Bernardo et al., 2003; Soch, Haynes, & Allefeld, 2016). From this perspective, models accounting for different temporal scales are seen as competing hypotheses, and can be systematically compared using Bayesian model comparison. In other words, modelling different temporal scales is justified if the multiscale model has a greater free-energy than a model without multiple scales. More generally, we recommend comparing the free-energy as a systematic way to arbitrate between different models, even when derived under different sets of assumptions. For instance, each temporal scale could be dynamically independent (e.g., frequency multiplexing); this would be considered as a hypothesis and mitigated against alternative hypotheses by comparing free-energies.

The impact of statistical models for multiscale dynamical systems in neuroscience is poised to grow. A remarkable example is the development of The Virtual Epileptic Patient, which helps construct personalized epilepsy models to aid in clinical comprehension of epileptic cases (Hashemi et al., 2020; V. Jirsa et al., 2023; V. K. Jirsa et al., 2017). Several promising directions for developing these models are worth exploring. First, HMMs of EEG and MEG signals could be extended to reflect the itinerancy between the various attractors of the (multi-stable) system arising from interconnected cortical population, for instance, using the escape rates introduced in (Cooray, Rosch, & Friston, 2023a, 2023b). Such transformation would allow one to relate the switching statistics of the Markov chain to properties of the network, effectively uncovering mechanistic explanations for state transitions. Second, employing models with continuous slow dynamics allows tracking synaptic gain in pyramidal neurons, believed to encode the precision (inverse variance) of top-down predictions on the lower levels of the predictive coding hierarchy. Investigating how these precisions change with experimental manipulations and over time would be particularly interesting. A third avenue for hierarchical models of multiscale time series would be to use theoretical frameworks like

structured flow on manifolds (SFM) to model empirical observations. For instance, applying an SFM-based generative model of neuronal activity to motor experiments, in the spirit of Churchland et al. (2012), Lara et al. (2018), and Saxena et al. (2022) could offer valuable insights into the dynamics of motor tasks.

In conclusion, by combining hierarchical state-space models with Bayesian analysis methods, time series data with different temporal scales can be linked and their interactions investigated. This has widespread application within neuroscience research, as well as in empirical science more generally.

ACKNOWLEDGMENTS

The authors are grateful to Robert Seymour for his useful suggestions.

AUTHOR CONTRIBUTIONS

Johan Medrano: Conceptualization; Visualization; Writing – original draft. Karl John Friston: Conceptualization; Writing – review & editing. Peter Zeidman: Conceptualization; Supervision; Writing – original draft.

FUNDING INFORMATION

Johan Medrano, Karl John Friston, Peter Zeidman, Wellcome Trust (<https://dx.doi.org/10.13039/100010269>), Award ID: 203147/Z/16/Z.

Q2 REFERENCES

- Adams, R. A., Bauer, M., Pinotsis, D., & Friston, K. J. (2016). Dynamic causal modelling of eye movements during pursuit: Confirming precision-encoding in V1 using MEG. *NeuroImage*, 132, 175–189. <https://doi.org/10.1016/j.neuroimage.2016.02.055>, PubMed: 26921713
- Arnold, L. (1995). *Random dynamical systems*. Berlin: Springer. <https://doi.org/10.1007/BFb0095238>
- Baker, A. P., Brookes, M. J., Rezek, I. A., Smith, S. M., Behrens, T., Smith, P. J. P., & Woolrich, M. (2014). Fast transient networks in spontaneous human brain activity. *Elife*, 3, e01867. <https://doi.org/10.7554/eLife.01867>, PubMed: 24668169
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron*, 76(4), 695–711. <https://doi.org/10.1016/j.neuron.2012.10.038>, PubMed: 23177956
- Bernardo, J., Bayarri, M., Berger, J., Dawid, A., Heckerman, D., Smith, A., ... others (2003). The variational Bayesian em algorithm for incomplete data: With application to scoring graphical model structures. *Bayesian Statistics*, 7, 453–464. <https://doi.org/10.1093/oso/9780198526155.003.0025>
- Bishop, C. M., & Nasrabadi, N. M. (2006). *Pattern recognition and machine learning* (Vol. 4) (No. 4). Springer.
- Boxler, P. (1989). A stochastic version of center manifold theory. *Probability Theory and Related Fields*, 83(4), 509–545. <https://doi.org/10.1007/BF01845701>
- Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nature Neuroscience*, 20(3), 340–352. <https://doi.org/10.1038/nn.4497>, PubMed: 28230845
- Cabral, J., Vidaurre, D., Marques, P., Magalhães, R., Silva Moreira, P., Miguel Soares, J., ... Kringelbach, M. L. (2017). Cognitive performance in healthy older adults relates to spontaneous switching between states of functional connectivity during rest. *Scientific Reports*, 7(1), 1–13. <https://doi.org/10.1038/s41598-017-05425-7>, PubMed: 28698644
- Carr, J. (1981). *Applications of centre manifold theory* (Vol. 35). Springer Science & Business Media. <https://doi.org/10.1007/978-1-4612-5929-9>
- Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A., & Fiete, I. (2019). The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9), 1512–1520. <https://doi.org/10.1038/s41593-019-0460-x>, PubMed: 31406365
- Chen, C., Kiebel, S. J., & Friston, K. J. (2008). Dynamic causal modelling of induced responses. *NeuroImage*, 41(4), 1293–1312. <https://doi.org/10.1016/j.neuroimage.2008.03.026>, PubMed: 18485744
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, 487(7405), 51–56. <https://doi.org/10.1038/nature11129>, PubMed: 22722855
- Cooray, G. K., Rosch, R. E., & Friston, K. J. (2023a). Global dynamics of neural mass models. *PLOS Computational Biology*, 19(2), e1010915. <https://doi.org/10.1371/journal.pcbi.1010915>, PubMed: 36763644
- Cooray, G. K., Rosch, R. E., & Friston, K. J. (2023b). Modelling cortical network dynamics. *arXiv preprint arXiv:2301.01141*. <https://doi.org/10.48550/arXiv.2301.01144>

- Courtiol, J., Guye, M., Bartolomei, F., Petkoski, S., & Jirsa, V. K. (2020). Dynamical mechanisms of interictal resting-state functional connectivity in epilepsy. *Journal of Neuroscience*, *40*(29), 5572–5588. <https://doi.org/10.1523/JNEUROSCI.0905-19.2020>, PubMed: 32513827
- D'Angelo, E., & Jirsa, V. (2022). The quest for multiscale brain modeling. *Trends in Neurosciences*, *45*(10), 777–790. <https://doi.org/10.1016/j.tins.2022.06.007>, PubMed: 35906100
- Daunizeau, J., Friston, K. J., & Kiebel, S. J. (2009). Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. *Physica D: Nonlinear Phenomena*, *238*(21), 2089–2118. <https://doi.org/10.1016/j.physd.2009.08.002>, PubMed: 19862351
- Deco, G., & Jirsa, V. K. (2012). Ongoing cortical activity at rest: Criticality, multistability, and ghost attractors. *Journal of Neuroscience*, *32*(10), 3366–3375. <https://doi.org/10.1523/JNEUROSCI.2523-11.2012>, PubMed: 22399758
- Demirtaş, M., Ponce-Alvarez, A., Gilson, M., Hagmann, P., Mantini, D., Betti, V., ... Deco, G. (2019). Distinct modes of functional connectivity induced by movie-watching. *NeuroImage*, *184*, 335–348. <https://doi.org/10.1016/j.neuroimage.2018.09.042>, PubMed: 30237036
- Depannemaecker, D., Destexhe, A., Jirsa, V., & Bernard, C. (2021). Modeling seizures: From single neurons to networks. *Seizure*, *90*, 4–8. <https://doi.org/10.1016/j.seizure.2021.06.015>, PubMed: 34219016
- Depannemaecker, D., Ezzati, A., Wang, H., Jirsa, V., & Bernard, C. (2023). From phenomenological to biophysical models of seizures. *Neurobiology of Disease*, *182*, 106131. <https://doi.org/10.1016/j.nbd.2023.106131>, PubMed: 37086755
- Echeverria-Altuna, I., Quinn, A. J., Zokaei, N., Woolrich, M. W., Nobre, A. C., & van Ede, F. (2022). Transient beta activity and cortico-muscular connectivity during sustained motor behaviour. *Progress in Neurobiology*, *214*, 102281. <https://doi.org/10.1016/j.pneurobio.2022.102281>, PubMed: 35550908
- Efron, B., & Morris, C. (1973). Stein's estimation rule and its competitors—An empirical Bayes approach. *Journal of the American Statistical Association*, *68*(341), 117–130. <https://doi.org/10.1080/01621459.1973.10481350>
- Ellis, G. F., Noble, D., & O'Connor, T. (2012). Top-down causation: An integrating theme within and across the sciences? *Interface Focus*, *2*, 1–3. <https://doi.org/10.1098/rsfs.2011.0110>
- Fliess, M., Lamnabhi, M., & Lamnabhi-Lagarrigue, F. (1983). An algebraic approach to nonlinear functional expansions. *IEEE Transactions on Circuits and Systems*, *30*(8), 554–570. <https://doi.org/10.1109/TCS.1983.1085397>
- Foias, C., Sell, G. R., & Temam, R. (1988). Inertial manifolds for non-linear evolutionary equations. *Journal of Differential Equations*, *73*(2), 309–353. [https://doi.org/10.1016/0022-0396\(88\)90110-6](https://doi.org/10.1016/0022-0396(88)90110-6)
- Fowlie, A. (2021). Neyman–Pearson lemma for Bayes factors. *Communications in Statistics-Theory and Methods*, *52*(15), 5379–5386. <https://doi.org/10.1080/03610926.2021.2007265>
- Fraser, A. M. (2008). *Hidden Markov models and dynamical systems*. SIAM. <https://doi.org/10.1137/1.9780898717747>
- Frässle, S., Aponte, E. A., Bollmann, S., Brodersen, K. H., Do, C. T., Harrison, O. K., ... Stephan, K. E. (2021). TAPAS: An open-source software package for translational neuromodeling and computational psychiatry. *Frontiers in Psychiatry*, *12*, 680811. <https://doi.org/10.3389/fpsy.2021.680811>, PubMed: 34149484
- Frässle, S., Harrison, S. J., Heinzle, J., Clementz, B. A., Tamminga, C. A., Sweeney, J. A., ... Stephan, K. E. (2021). Regression dynamic causal modeling for resting-state fMRI. *Human Brain Mapping*, *42*(7), 2159–2180. <https://doi.org/10.1002/hbm.25357>, PubMed: 33539625
- Frässle, S., Lomakina, E. I., Kasper, L., Manjaly, Z. M., Leff, A., Pruessmann, K. P., ... Stephan, K. E. (2018). A generative model of whole-brain effective connectivity. *NeuroImage*, *179*, 505–529. <https://doi.org/10.1016/j.neuroimage.2018.05.058>, PubMed: 29807151
- Friston, K. J., Flandin, G., & Razi, A. (2022). Dynamic causal modelling of COVID-19 and its mitigations. *Scientific Reports*, *12*(1), 12419. <https://doi.org/10.1038/s41598-022-16799-8>, PubMed: 35859054
- Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *NeuroImage*, *19*(4), 1273–1302. [https://doi.org/10.1016/S1053-8119\(03\)00202-7](https://doi.org/10.1016/S1053-8119(03)00202-7), PubMed: 12948688
- Friston, K. J., Li, B., Daunizeau, J., & Stephan, K. E. (2011). Network discovery with DCM. *NeuroImage*, *56*(3), 1202–1221. <https://doi.org/10.1016/j.neuroimage.2010.12.039>, PubMed: 21182971
- Friston, K. J., Litvak, V., Oswal, A., Razi, A., Stephan, K. E., Wijk, B. C. V., ... Zeidman, P. (2016). Bayesian model reduction and empirical bayes for group (DCM) studies. *NeuroImage*, *128*, 413–431. <https://doi.org/10.1016/j.neuroimage.2015.11.015>, PubMed: 26569570
- Friston, K. J., Mechelli, A., Turner, R., & Price, C. J. (2000). Nonlinear responses in fMRI: The balloon model, volterra kernels, and other hemodynamics. *NeuroImage*, *12*(4), 466–477. <https://doi.org/10.1006/nimg.2000.0630>, PubMed: 10988040
- Friston, K. J., Parr, T., Zeidman, P., Razi, A., Flandin, G., Daunizeau, J., ... Lambert, C. (2020). Second waves, social distancing, and the spread of COVID-19 across the usa. *Wellcome Open Research*, *5*. <https://doi.org/10.12688/wellcomeopenres.15986.1>
- Friston, K. J., Redish, A. D., & Gordon, J. A. (2017). Computational nosology and precision psychiatry. *Computational Psychiatry*, *1*, 2–23. https://doi.org/10.1162/cpsy_a_00001, PubMed: 29400354
- Gallego, J. A., Perich, M. G., Naufel, S. N., Ethier, C., Solla, S. A., & Miller, L. E. (2018). Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications*, *9*(1), 4233. <https://doi.org/10.1038/s41467-018-06560-z>, PubMed: 30315158
- Graham, A. M., Marr, M., Buss, C., Sullivan, E. L., & Fair, D. A. (2021). Understanding vulnerability and adaptation in early brain development using network neuroscience. *Trends in Neurosciences*, *44*, 276–288. <https://doi.org/10.1016/j.tins.2021.01.008>, PubMed: 33663814
- Guckenheimer, J., & Holmes, P. (1983). *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields* (Vol. 42). Springer Science & Business Media. <https://doi.org/10.1007/978-1-4612-1140-2>
- Haken, H. (1977). Synergetics. *Physics Bulletin*, *28*(9), 412. <https://doi.org/10.1088/0031-9112/28/9/027>
- Haken, H. (1978). *Synergetics: An introduction: Nonequilibrium phase transitions and self-organization in physics, chemistry,*

- and biology (2nd edn.). Berlin: Springer-Verlag. <https://doi.org/10.1007/978-3-642-96469-5>
- Haken, H., Kelso, J. S., & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, 51(5), 347–356. <https://doi.org/10.1007/BF00336922>, PubMed: 3978150
- Hashemi, M., Vattikonda, A., Sip, V., Guye, M., Bartolomei, F., Woodman, M. M., & Jirsa, V. K. (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*, 217, 116839. <https://doi.org/10.1016/j.neuroimage.2020.116839>, PubMed: 32387625
- Heitmann, S., & Breakspear, M. (2018). Putting the “dynamic” back into dynamic functional connectivity. *Network Neuroscience*, 2(2), 150–174. https://doi.org/10.1162/netn_a_00041, PubMed: 30215031
- Higgins, C., Liu, Y., Vidaurre, D., Kurth-Nelson, Z., Dolan, R., Behrens, T., & Woolrich, M. (2021). Replay bursts in humans coincide with activation of the default mode and parietal alpha networks. *Neuron*, 109(5), 882–893. <https://doi.org/10.1016/j.neuron.2020.12.007>, PubMed: 33357412
- Jafarian, A., Zeidman, P., Wykes, R. C., Walker, M., & Friston, K. J. (2021). Adiabatic dynamic causal modelling. *NeuroImage*, 238, 118243. <https://doi.org/10.1016/j.neuroimage.2021.118243>, PubMed: 34116151
- Jirsa, V. (2020). Structured flows on manifolds as guiding concepts in brain science. In *Selbstorganisation—Ein Paradigma für die Humanwissenschaften* (pp. 89–102). Springer. https://doi.org/10.1007/978-3-658-29906-4_6
- Jirsa, V., & Sheheitli, H. (2022). Entropy, free energy, symmetry and dynamics in the brain. *Journal of Physics: Complexity*, 3(1), 015007. <https://doi.org/10.1088/2632-072X/ac4bec>
- Jirsa, V., Wang, H., Triebkorn, P., Hashemi, M., Jha, J., Gonzalez-Martinez, J., ... Bartolomei, F. (2023). Personalised virtual brain models in epilepsy. *The Lancet Neurology*, 22(5), 443–454. [https://doi.org/10.1016/S1474-4422\(23\)00008-X](https://doi.org/10.1016/S1474-4422(23)00008-X), PubMed: 36972720
- Jirsa, V. K., McIntosh, A. R., & Huys, R. (2019). Grand unified theories of the brain need better understanding of behavior: The two-tiered emergence of function. *Ecological Psychology*, 31(3), 152–165. <https://doi.org/10.1080/10407413.2019.1615207>
- Jirsa, V. K., Proix, T., Perdikis, D., Woodman, M. M., Wang, H., Gonzalez-Martinez, J., ... Bartolomei, F. (2017). The virtual epileptic patient: Individualized whole-brain models of epilepsy spread. *NeuroImage*, 145, 377–388. <https://doi.org/10.1016/j.neuroimage.2016.04.049>, PubMed: 27477535
- 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. <https://doi.org/10.1093/brain/awu133>, PubMed: 24919973
- Kalman, R. E. (1960). On the general theory of control systems. In *Proceedings first international conference on automatic control, Moscow, USSR* (pp. 481–492). [https://doi.org/10.1016/S1474-6670\(17\)70094-8](https://doi.org/10.1016/S1474-6670(17)70094-8)
- Kanai, R., Komura, Y., Shipp, S., & Friston, K. (2015). Cerebral hierarchies: Predictive processing, precision and the pulvinar. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1668), 20140169. <https://doi.org/10.1098/rstb.2014.0169>, PubMed: 25823866
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90, 773–795. <https://doi.org/10.1080/01621459.1995.10476572>
- Kass, R. E., & Steffey, D. (1989). Approximate Bayesian inference in conditionally independent hierarchical models (parametric empirical bayes models). *Journal of the American Statistical Association*, 84(407), 717–726. <https://doi.org/10.1080/01621459.1989.10478825>
- Kelso, J. S. (2012). Multistability and metastability: Understanding dynamic coordination in the brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1591), 906–918. <https://doi.org/10.1098/rstb.2011.0351>, PubMed: 22371613
- Knobloch, E., & Wiesenfeld, K. (1983). Bifurcations in fluctuating systems: The center-manifold approach. *Journal of Statistical Physics*, 33(3), 611–637. <https://doi.org/10.1007/BF01018837>
- Kuang, L., Han, X., Chen, K., Caselli, R. J., Reiman, E. M., Wang, Y., & Alzheimer’s Disease Neuroimaging Initiative. (2019). A concise and persistent feature to study brain resting-state network dynamics: Findings from the Alzheimer’s disease neuroimaging initiative. *Human Brain Mapping*, 40(4), 1062–1081. <https://doi.org/10.1002/hbm.24383>, PubMed: 30569583
- Kuramoto, Y., & Nakao, H. (2019). On the concept of dynamical reduction. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 377(2160), 1–25. <https://doi.org/10.1098/rsta.2019.0041>, PubMed: 31656146
- Langdon, C., Genkin, M., & Engel, T. A. (2023). A unifying perspective on neural manifolds and circuits for cognition. *Nature Reviews Neuroscience*, 1–15. <https://doi.org/10.1038/s41583-023-00693-x>, PubMed: 37055616
- Lara, A. H., Cunningham, J. P., & Churchland, M. M. (2018). Different population dynamics in the supplementary motor area and motor cortex during reaching. *Nature Communications*, 9(1), 2754. <https://doi.org/10.1038/s41467-018-05146-z>, PubMed: 30013188
- Lee, C. S., Aly, M., & Baldassano, C. (2021). Anticipation of temporally structured events in the brain. *Elife*, 10, e64972. <https://doi.org/10.7554/eLife.64972>, PubMed: 33884953
- Lisi, G., Rivela, D., Takai, A., & Morimoto, J. (2018). Markov switching model for quick detection of event related desynchronization in EEG. *Frontiers in Neuroscience*, 12, 24. <https://doi.org/10.3389/fnins.2018.00024>, PubMed: 29449799
- Lopez-Sola, E., Sanchez-Todo, R., Lleal, È., Köksal-Ersöz, E., Yochum, M., Makhlova, J., ... Ruffini, G. (2022). A personalizable autonomous neural mass model of epileptic seizures. *Journal of Neural Engineering*, 19(5), 055002. <https://doi.org/10.1088/1741-2552/ac8ba8>, PubMed: 35995031
- McIntosh, A. R., & Jirsa, V. K. (2019). The hidden repertoire of brain dynamics and dysfunction. *Network Neuroscience*, 3(4), 994–1008. https://doi.org/10.1162/netn_a_00107, PubMed: 31637335
- Meer, J. N., Breakspear, M., Chang, L. J., Sonkusare, S., & Cocchi, L. (2020). Movie viewing elicits rich and reliable brain state dynamics. *Nature Communications*, 11(1), 1–14. <https://doi.org/10.1038/s41467-020-18717-w>, PubMed: 33020473

- Moran, R., Pinotsis, D. A., & Friston, K. (2013). Neural masses and fields in dynamic causal modeling. *Frontiers in Computational Neuroscience*, 7, 57. <https://doi.org/10.3389/fncom.2013.00057>, PubMed: 23755005
- Moran, R. J., Kiebel, S. J., Stephan, K. E., Reilly, R. B., Daunizeau, J., & Friston, K. J. (2007). A neural mass model of spectral responses in electrophysiology. *NeuroImage*, 37, 706–720. <https://doi.org/10.1016/j.neuroimage.2007.05.032>, PubMed: 17632015
- 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(4), 1694–1708. <https://doi.org/10.1016/j.neuroimage.2011.01.012>, PubMed: 21238593
- Neyman, J., & Pearson, E. S. (1933). IX. on the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694–706), 289–337. <https://doi.org/10.1098/rsta.1933.0009>
- Núñez, P., Poza, J., Gómez, C., Rodríguez-González, V., Hillebrand, A., Tewarie, P., ... Hornero, R. (2021). Abnormal meta-state activation of dynamic brain networks across the Alzheimer spectrum. *NeuroImage*, 232, 117898. <https://doi.org/10.1016/j.neuroimage.2021.117898>, PubMed: 33621696
- O'Neill, G. C., Tewarie, P., Vidaurre, D., Liuzzi, L., Woolrich, M. W., & Brookes, M. J. (2018). Dynamics of large-scale electrophysiological networks: A technical review. *NeuroImage*, 180, 559–576. <https://doi.org/10.1016/j.neuroimage.2017.10.003>, PubMed: 28988134
- 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.
- Pillai, A. S., & Jirsa, V. K. (2017). Symmetry breaking in space-time hierarchies shapes brain dynamics and behavior. *Neuron*, 94(5), 1010–1026. <https://doi.org/10.1016/j.neuron.2017.05.013>, PubMed: 28595045
- Ponce-Alvarez, A., Deco, G., Hagmann, P., Romani, G. L., Mantini, D., & Corbetta, M. (2015). Resting-state temporal synchronization networks emerge from connectivity topology and heterogeneity. *PLoS Computational Biology*, 11(2), e1004100. <https://doi.org/10.1371/journal.pcbi.1004100>, PubMed: 25692996
- Ponce-Alvarez, A., Kringelbach, M. L., & Deco, G. (2023). Critical scaling of whole-brain resting-state dynamics. *Communications Biology*, 6(1), 627. <https://doi.org/10.1038/s42003-023-05001-y>, PubMed: 37301936
- Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), 257–286. <https://doi.org/10.1109/5.18626>
- Rabinovich, M. I., Huerta, R., Varona, P., & Afraimovich, V. S. (2008). Transient cognitive dynamics, metastability, and decision making. *PLoS Computational Biology*, 4(5), e1000072. <https://doi.org/10.1371/journal.pcbi.1000072>, PubMed: 18452000
- Rabinovich, M. I., & Varona, P. (2018). Discrete sequential information coding: Heteroclinic cognitive dynamics. *Frontiers in Computational Neuroscience*, 12, 73. <https://doi.org/10.3389/fncom.2018.00073>, PubMed: 30245621
- Roberts, A. J. (2014). *Model emergent dynamics in complex systems* (Vol. 20). SIAM. <https://doi.org/10.1137/1.9781611973563>
- Rosch, R., Baldeweg, T., Moeller, F., & Baier, G. (2018). Network dynamics in the healthy and epileptic developing brain. *Network Neuroscience*, 2(1), 41–59. https://doi.org/10.1162/NETN_a_00026, PubMed: 29911676
- Rosch, R. E., Hunter, P. R., Baldeweg, T., Friston, K. J., & Meyer, M. P. (2018). Calcium imaging and dynamic causal modelling reveal brain-wide changes in effective connectivity and synaptic dynamics during epileptic seizures. *PLoS Computational Biology*, 14, e1006375. <https://doi.org/10.1371/journal.pcbi.1006375>, PubMed: 30138336
- Rosch, R. E., Wright, S., Cooray, G., Papadopoulou, M., Goyal, S., Lim, M., ... Friston, K. J. (2018). NMDA-receptor antibodies alter cortical microcircuit dynamics. *Proceedings of the National Academy of Sciences*, 115(42), E9916–E9925. <https://doi.org/10.1073/pnas.1804846115>, PubMed: 30262654
- Sanz Leon, P., Knock, S. A., Woodman, M. M., Domide, L., Mersmann, J., McIntosh, A. R., & Jirsa, V. (2013). The virtual brain: A simulator of primate brain network dynamics. *Frontiers in Neuroinformatics*, 7, 10. <https://doi.org/10.3389/fninf.2013.00010>, PubMed: 23781198
- Saxena, S., Russo, A. A., Cunningham, J., & Churchland, M. M. (2022). Motor cortex activity across movement speeds is predicted by network-level strategies for generating muscle activity. *Elife*, 11, e67620. <https://doi.org/10.7554/eLife.67620>, PubMed: 35621264
- Shain, C., Blank, I. A., van Schijndel, M., Schuler, W., & Fedorenko, E. (2020). fMRI reveals language-specific predictive coding during naturalistic sentence comprehension. *Neuropsychologia*, 138, 107307. <https://doi.org/10.1016/j.neuropsychologia.2019.107307>, PubMed: 31874149
- Smailovic, U., Koenig, T., Laukka, E. J., Kalpouzos, G., Andersson, T., Winblad, B., & Jelic, V. (2019). EEG time signature in Alzheimer's disease: Functional brain networks falling apart. *NeuroImage: Clinical*, 24, 102046. <https://doi.org/10.1016/j.nicl.2019.102046>, PubMed: 31795039
- Soch, J., Haynes, J.-D., & Allefeld, C. (2016). How to avoid misspecification in GLM-based fMRI data analysis: Cross-validated Bayesian model selection. *NeuroImage*, 141, 469–489. <https://doi.org/10.1016/j.neuroimage.2016.07.047>, PubMed: 27477536
- Stephan, K. E., Kasper, L., Harrison, L. M., Daunizeau, J., Ouden, E. M. D., Breakspear, M., & Friston, K. J. (2008). Nonlinear dynamic causal models for fMRI. *NeuroImage*, 42(2), 649–662. <https://doi.org/10.1016/j.neuroimage.2008.04.262>, PubMed: 18565765
- Symmonds, M., Moran, C. H., Leite, M. I., Buckley, C., Irani, S. R., Stephan, K. E., ... Moran, R. J. (2018). Ion channels in EEG: Isolating channel dysfunction in nmda receptor antibody encephalitis. *Brain*, 141(6), 1691–1702. <https://doi.org/10.1093/brain/awy107>, PubMed: 29718139
- Tognoli, E., & Kelso, J. S. (2014). The metastable brain. *Neuron*, 81(1), 35–48. <https://doi.org/10.1016/j.neuron.2013.12.022>, PubMed: 24411730
- Vidaurre, D., Abeysuriya, R., Becker, R., Quinn, A. J., Alfaro-Almagro, F., Smith, S. M., & Woolrich, M. W. (2018). Discovering dynamic brain networks from big data in rest and task. *NeuroImage*, 180, 646–656. <https://doi.org/10.1016/j.neuroimage.2017.06.077>, PubMed: 28669905

Vidaurre, D., Quinn, A. J., Baker, A. P., Dupret, D., Tejero-Cantero, A., & Woolrich, M. W. (2016). Spectrally resolved fast transient brain states in electrophysiological data. *NeuroImage*, *126*, 81–95. <https://doi.org/10.1016/j.neuroimage.2015.11.047>, PubMed: 26631815

Vidaurre, D., Smith, S. M., & Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, *114*(48), 12827–12832. <https://doi.org/10.1073/pnas.1705120114>, PubMed: 29087305

Wang, X., Huang, W., Su, L., Xing, Y., Jessen, F., Sun, Y., ... Han, Y. (2020). Neuroimaging advances regarding subjective cognitive decline in preclinical Alzheimer's disease. *Molecular Neurodegeneration*, *15*(1), 1–27. <https://doi.org/10.1186/s13024-020-00395-3>, PubMed: 32962744

Zarghami, T. S., & Friston, K. J. (2020). Dynamic effective connectivity. *NeuroImage*, *207*, 116453. <https://doi.org/10.1016/j.neuroimage.2019.116453>, PubMed: 31821868

Uncorrected Proof