NeuroImage 44 (2009) 701–714

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Population dynamics under the Laplace assumption ☆

André C. Marreiros *, Stefan J. Kiebel, Jean Daunizeau, Lee M. Harrison, Karl J. Friston

The Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London, WC1N 3BG, UK

ARTICLE INFO

Article history:
Received 4 August 2008
Revised 30 September 2008
Accepted 10 October 2008
Available online 25 October 2008

Keywords:
Neural-mass models
Nonlinear
Modelling
Laplace assumption
Mean-field
Neuronal

ABSTRACT

In this paper, we describe a generic approach to modelling dynamics in neuronal populations. This approach models a full density on the states of neuronal populations but finesses this high-dimensional problem by re-formulating density dynamics in terms of ordinary differential equations on the sufficient statistics of the densities considered (c.f., the method of moments). The particular form for the population density we adopt is a Gaussian density (c.f., the Laplace assumption). This means population dynamics are described by equations governing the evolution of the population's mean and covariance. We derive these equations from the Fokker-Planck formalism and illustrate their application to a conductance-based model of neuronal exchanges. One interesting aspect of this formulation is that we can uncouple the mean and covariance to furnish a neural-mass model, which rests only on the populations mean. This enables us to compare equivalent mean-field and neural-mass models of the same populations and evaluate, quantitatively, the contribution of population variance to the expected dynamics. The mean-field model presented here will form the basis of a dynamic causal model of observed electromagnetic signals in future work.

Introduction

Mean-field models of neuronal dynamics have a long history, spanning a half-century (e.g., Beurle 1956). Models are essential for neuroscience, in the sense that most interesting questions about the brain pertain to neuronal mechanisms and processes that are not directly observable (Tass 2003; Breakspear et al., 2006). This means that questions about neuronal function are generally addressed by inference on models or their parameters; where the model links hidden neuronal processes to our observations and questions (Valdes et al., 1999). Broadly speaking, models are used to generate data to study emergent behaviours. Alternatively, they can be used as forward or observation models (e.g., dynamic causal models), which are inverted given empirical data (David et al., 2006, Kiebel et al., 2006). This inversion allows one to select the best model (i.e., hypothesis), given some data and make probabilistic statements about the parameters of that model (e.g., Penny et al., 2004).

In particular, mean-field models are appropriate for data that reflect the behaviour of neuronal populations, such as the electro-encephalogram (EEG), magnetoencephalogram (MEG) and functional magnetic resonance imaging (fMRI) data. The most prevalent models of neuronal populations or ensembles are based upon the so-called *mean-field approximation*. This approximation replaces the time-averaged discharge rate of individual neurons with a common time-dependent population activity (ensemble average; Knight 2000; Haskell et al., 2001). The mean-field approximation is used extensively in statistical physics for otherwise computationally or analytically intractable problems. An exemplary approach, owing to Boltzmann and Maxwell, is the approximation of the motion of molecules in a gas by mean-field terms such as temperature and pressure. Similarly, evoked response potentials (ERPs) represent the average response over millions of neurons, where the mean-field approximation describes the time-dependent distribution of the average population response. This is possible because the dynamics of the mean of the density are much less stochastic than the response of a single neuron. This makes it feasible to develop algorithms that use Bayesian inference to infer neuronal parameters given measured responses, using mean-field models (e.g., Harrison et al., 2005).

Usually, neural-mass models are used to model the evolution of the mean response or the response at steady state. Mean-field approximations go further and model the full distribution of the population response. However, mean-field models can be computationally expensive, because one has to consider the density at all points in neuronal state-space as opposed to a single quantity (e.g., the mean). In this paper, we present an approach that simplifies the mean-field model by using the Laplace approximation: Under the Laplace approximation, the population or ensemble density assumes a Gaussian form, whose sufficient statistics comprise the conditional mean and covariance. In contrast to neural-mass models, this allows one to model interactions between the first two moments (i.e., mean and variance) of neuronal states. In a subsequent paper, we will use the Laplace and neural mass approximations presented here as generative

☆ Software Note. Matlab demonstration and modelling routines referred to in this paper are available as academic freeware as part of the SPM software from http://www.fil.ion.ucl.ac.uk/spm (neural models toolbox).

* Corresponding author. Fax: +44 207 813 1445.
E-mail address: a.marreiros@fil.ion.ucl.ac.uk (A.C. Marreiros).

1053-8119/$ – see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2008.10.008

models of electrophysiological responses to sensory input. This subsequent paper will use Bayesian model comparison to compare both models and establish whether empirical responses contain evidence for a role of variance in shaping population dynamics.

The Laplace approximation is a ubiquitous device in statistical physics and machine learning and finesses difficult integration problems when integrating over probability densities (see Friston et al 2007; Chumbley et al 2007). Exactly the same device is used here to furnish a simple scheme for modeling density dynamics. Because the sufficient statistics of a Gaussian density can be specified in terms of the first two moments, the ensuing scheme is formally identical to the second-moment method described by Rodriguez and Tuckwell (1996). The method of moments (Rodriguez and Tuckwell 1996, 1998), replaces a system of stochastic differential equations (describing the states of an ensemble) with deterministic equations describing the evolution of the sufficient statistics or moments of an ensemble density. This approach was first applied to a FitzHugh-Nagumo (FN) neuron (Rodriguez and Tuckwell 1996, Tuckwell and Rodriguez 1998) and later to Hodgkin Huxley (HH) neurons (Rodriguez and Tuckwell 1998, 2000). This approach assumes that the distributions of the variables are approximately Gaussian so that they can be characterized by their first and second order moments; i.e., the means and covariances. In related work, Hasegawa described a dynamical mean-field approximation (DMA) to simulate the activities of a neuronal network. This method allows for qualitative or semi-quantitative inference on the properties of ensembles or clusters of FN and HH neurons; see Hasegawa (2003a,b).

This paper comprises three sections. In the first, we provide the background to modelling neuronal dynamics with mean-field and neural-mass models. In the second section, we derive a generic mean-field treatment of neuronal dynamics starting with any equations of motion. This treatment is based on a Laplace approximation to the ensemble density, and is formulated compactly, in terms of the equations of motion for the sufficient statistics of the ensemble density. This approach reduces to a neural-mass model when the second-order statistics (i.e., variance) of neuronal states are ignored. We will illustrate how neuronal state equations are reformulated as a mean-field approximation, using a simple conductance-based model (c.f., Morris and Lecar, 1981). In the third section, we establish the validity of the Laplace approximation by comparing the response of simulated ensembles of neurons to responses under the Laplace and neural-mass assumptions. The key behaviour we are interested in is the coupling between the mean and variance of the ensemble, which is lost in the neural-mass approximations.

Mean field and neural-masses

What follows is a brief summary of the material in Deco et al (in press), which provides a full account of mean-field models in neuroscience. The most prevalent models of neuronal populations or ensembles are based on the mean-field approximation. The basic idea behind these models is to approximate a very high dimensional probability distribution with the product of a number of simpler (marginal) densities. Its utility is best seen in the context of ensemble or population density models.

Mean-field models

Ensemble models attempt to model the dynamics of large populations of neurons. Any single neuron can have a number of attributes; for example, post-synaptic membrane depolarisation, V, capacitive current I or the time since the last action potential, T. Each attribute induces a dimension in the state or *phase-space* of a neuron. In this example, the phase-space would be three-dimensional and the state of each neuron would correspond to a point $x=\{V,I,T\}\in R^3$ or particle in phase-space. Imagine a very large number of such neurons that populate phase-space with a density, $q(x,t)$. As the state of each neuron evolves, the points will flow through phase-space and the *ensemble density*; $q(x,t)$ will evolve until it reaches some steady-state or equilibrium. It is the evolution of the density *per se* that is characterised in ensemble density methods. These models are particularly attractive because the density dynamics conform to a simple equation; *the Fokker Planck equation* (Risken 1996; Dayan and Abbott 2001; Frank et al., 2001; Gerstner and Kistler 2002)

$$\dot{q} = -\nabla\cdot fq + \nabla\cdot D\nabla q = -\sum_{i=1}^{n}\frac{\partial(f_i q)}{\partial x_i} + \sum_{i,j=1}^{n}\left(\frac{\partial}{\partial x_i}D_{ij}\frac{\partial}{\partial x_j}\right)q \tag{1}$$

For n states; $x\in\Re^n$. The equation comprises flow and dispersion terms, which embed the assumptions about the dynamics and random fluctuations. The flow, $f(x,t)$ and dispersion, $D(x,t)$ constitute our model at the neuronal level. This level of description is usually framed as a *stochastic differential equation* (SDE) that describes how the states evolve as functions of each other and some random fluctuations

$$dx = f(x)dt + \sigma dw. \tag{2}$$

Where, $D=\frac{1}{2}\sigma^2$ and $w(t)$ is a standard *Wiener process* (where, in one dimension $w(t+\Delta t)-w(t)\sim N(0,\Delta t)$). Under the Fokker-Planck formalism, even if the dynamics of each neuron are very complicated, or indeed chaotic, the density dynamics remain simple, linear and deterministic. In short, for any model of neuronal dynamics, specified as a stochastic differential equation, there is a deterministic linear equation that can be integrated to generate ensemble dynamics. However, there is a problem; the dimensionality of phase-space can become unmanageably large, if we consider too many neuronal states or different types of neuron. Generally speaking, full ensemble models of realistic systems are computationally intractable. However, we can use a mean-field approximation to finesse this problem.

The mean-field approximation

Consider the states of m sorts of neuron, each with n states; then the states $x=x^{(1)},\ldots,x^{(m)}\in\Re^{n\times m}$ could have a large $n\times m$ dimensionality. However, if we assume the density factorises over the m populations

$$q(x)\approx\prod_{i=1}^{m} q\left(x^{(i)}\right) \tag{3}$$

we have only to deal with n-dimensional states $x^{(i)}\in\Re^n$. However, by factorising the density into marginal densities we have effectively assumed that they are independent. This implausible assumption can be circumvented by coupling the ensembles so that the flow in the phase-space of one ensemble, $f(x^{(i)},\mu)$ depends upon the others; through mean-field quantities $\mu^{(i)}=\mu(q(x^{(i)}))$. These are *phase-functions* of the ensemble densities. These mean-field effects could come from the same ensemble and model interactions among neurons in the same population. The ensuing dynamics conform to a series of coupled nonlinear Fokker-Planck equations (Frank 2004). Typically, these phase-functions return the average state (e.g., mean depolarisation or firing). It is important to realise that coupling ensembles through mean-field quantities, $\mu\in\Re^n$ entails strong assumptions about the nature of the interactions: specifically, the dynamics or fluctuations in one member of an ensemble cannot affect a member of another ensemble. Instead, all the neurons in one ensemble are affected identically by the average behaviour of another ensemble. In many instances, this is a reasonable approximation but, clearly, it makes the exact form of the mean-field approximation an important consideration. In subsequent work, we will incorporate the mean-field model of this paper into dynamic causal models of distributed neuronal sources. In this context, the coupling above determines how one neuronal source influences another; i.e., it corresponds to effective connectivity (Friston et al., 2003; David et al., 2006).

Even for a single ensemble the dimensionality of $\in\Re^n$ may preclude numerical or analytic analysis. One can simplify the model by summarising the dynamics with a small number of states. In the limit, one can reduce the dynamics to a single neuronal state $x\in\Re$. An important example is when the state is voltage, i.e., $x=\{V\}$. For example, Gerstner and Kistler (2002) formulate the dynamics of an ensemble of *leaky integrate and fire* neurons with equations of motion

$$f(V) = -\frac{g_L}{C}(V-V_L) + \mu \tag{4}$$

using the Fokker-Planck equation (Eq. (1)), with boundary conditions on $q(V, t)$ that model spiking and a re-setting of the membrane potential. Here, C represents membrane capacitance and g_L a leakage conductance An alternative method is to use the auxiliary variable T (time elapsed since last spike) to parameterize the *refractory density*, $q(T, t)$; see Eggert and van Hemmen (2001). Chizhov et al., (2006) have refined this approach to account for fast and slow ionic currents, with some compelling results.

In summary, one can approximate an ensemble density on a high-dimensional phase-space with a series of low-dimension ensembles that are coupled through mean-field effects. The product of these marginal densities is then used to approximate the full density. Critically, the mean-field coupling induces nonlinear dependencies among the density dynamics of each ensemble. This typically requires a nonlinear Fokker-Planck equation for each ensemble. The Fokker-Planck equation prescribes the evolution of the ensemble dynamics, given any initial conditions and equations of motion that constitute our neuronal model. However, it does not specify how to encode or parameterize the density. There are several approaches to density parameterization (Knight 2000; Nykamp and Tranchina 2000; Omurtag et al., 2000; Haskell et al., 2001; Casti et al., 2002; Sirovich 2003). These include binning the phase-space and using a discrete approximation to a free-form density. However, this can lead to a vast number of differential equations, especially if there are multiple states for each population. One solution to this is to reduce the dimension of the phase-space to render the integration of the Fokker-Planck more tractable (e.g., Chizhov and Graham 2007). Alternatively, one can assume the density has a fixed parametric form and deal only with its sufficient statistics (Rodriguez and Tuckwell 1996, 1998; Hasegawa 2003a,b). The simplest form is a delta-function or point mass; under this assumption we get neural-mass models.

Neural-mass models

Neural-mass models can be regarded as a special case of ensemble density models, where we summarize the ensemble density with a single number. Early examples can be found in the work of Beurle (1956) and Griffith (1963, 1965). The term *mass action model* was coined by Freeman (1975) as an alternative to density dynamics. Assuming that the equilibrium density has a *point mass* (i.e., a delta function), we can motivate the description above in terms of the expected value of the state, μ; under the assumption that the equilibrium density has a *point mass* (i.e., a delta function). This is one perspective on why these simple mean-field models are called *neural-mass models*. In short, we replace the full ensemble density with a mass at a particular point and then summarize the density dynamics by the location of that mass. What we are left with is a set of non-linear differential equations describing the dynamic evolution of this mode. But what have we thrown away? In the full nonlinear Fokker-Planck formulation, different phase-functions or probability density moments could couple to each other; both within and between ensembles. For example, the average depolarisation in one ensemble could be effected by the dispersion or variance of depolarisation in another. See Deco et al (in press) for a nice example from the work of Michael Breakspear. In neural-mass models, one ignores this potential dependency because only the expectations or first moments are coupled. There are several devices that are used to compensate for this simplification. Perhaps the most ubiquitous is the use of a sigmoid function $\varsigma(V)$ relating expected depolarisation to expected firing-rate (Freeman 1975; Marreiros et al., 2008): This implicitly encodes variability in the post-synaptic depolarisation, relative to the potential at which the neuron would fire. A common form for neural-mass equations of motion posits a second-order differential equation for expected voltage μ_V or, equivalently, two coupled first-order equations, where

$$\begin{aligned}\left(\frac{\partial^2}{\partial t^2} + 2\kappa\frac{\partial}{\partial t} + \kappa^2\right)\mu_V = \kappa^2\gamma_\varsigma(\mu_V) \Rightarrow \\ \dot{\mu}_V = \mu_I \\ \dot{\mu}_I = \kappa^2\gamma_\varsigma(\mu_V) - 2\kappa\mu_I - \kappa^2\mu_V\end{aligned} \tag{5}$$

where m_I can be regarded as capacitive current. The input $gV(m_V)$ is commonly construed as firing-rate (or pulse-density) and is a sigmoid function of mean voltage of the same of another ensemble. The coupling constant g scales the amplitude of this mean-field effect. The constant κ controls the rise and decay of the implicit (synaptic) impulse response $G(t)$ to input; convolving input with this impulse response kernel gives the expected voltage

$$\begin{aligned}\mu_V(t) &= \int G(t-t')\varsigma(\mu_V(t))dt' \\ G(t) &= \gamma\kappa^2 t\exp(-\kappa t)\end{aligned} \tag{6}$$

This form of neural-mass model has been used extensively to model electrophysiological recordings (e.g., Lopes da Silva et al., 1974; Zetterberg et al., 1978; Elbert et al., 1994; Jansen and Rit 1995; Kincses et al., 1999; Wendling et al., 2000; David and Friston 2003; Moran et al., 2007) and has been used recently as the basis of a generative model for event-related potentials that can be inverted using real data (Valdes et al., 1999; Jansen et al., 2001; Friston et al., 2003; Kiebel et al., 2006; David et al., 2006; Moran et al 2007, 2008).

In short, neural-mass models are special cases of ensemble density models that are furnished by ignoring all but the expectation or mode of the ensemble density. This affords the considerable simplification of the dynamics and allows one to focus on the behaviour of a large number of ensembles, without having to worry about an explosion in the number of dimensions or differential equations one has to integrate. An important generalisation of neural-mass models, which allow for states that are functionals of position on the cortical sheet, are referred to as *neural-field models* (Jirsa and Haken 1996; Wright et al., 2003; Robinson et al., 2003; Breakspear et al., 2006). Deco et al (in press) provide a comprehensive overview of neural-mass and field models, to which the interested reader is referred.

Summary

In conclusion, statistical descriptions of neuronal ensembles can be formulated in terms of a Fokker-Planck equation; an equation prescribing the evolution of a probability density on some phase-space. The high dimensionality and complexity of these Fokker-Planck formalisms can be finessed with a mean-field approximation to give nonlinear Fokker-Planck equations, describing the evolution of separable ensembles that are coupled by mean-field effects. By parameterising the densities in terms of their sufficient statistics, these partial differential equations can be reduced to ordinary differential equations describing the evolution of the statistics. In the simplest case, we can use a single statistic corresponding to the expectation or mode of the probability for each ensemble. This can be regarded as encoding the location of a probability mass. In what follows, we consider what would happen if the sufficient statistics included both the mean and dispersion.

Ensemble dynamics under the laplace assumption

In this section, we derive a general mean-field reduction for neural dynamics formulated with any set of ordinary differential

equations. This is formally equivalent to the method of moments (MM) proposed by Rodriguez and Tuckwell (1996, 1998) for summarising density dynamics. In the next section, we apply the treatment to the equations used in dynamic causal modelling (DCM) of electrophysiological responses. The treatment here rests on summarising the ensemble density with a fixed form; namely, a Gaussian density. This corresponds to the Laplace assumption made in mean-field treatments in variational or ensemble learning in statistics. Here we use this approach to reduce a very high-dimensional integration problem into the manageable integration of the sufficient statistics (e.g., moments) of the ensemble density. The sufficient statistics are those quantities needed to define a particular density, in this case the mean $\mu^{(i)}$ and covariance $\Sigma^{(i)}$ of the states of the i-th population, with a multivariate normal distribution; $q(x^{(i)}) = N(\mu^{(i)}, \Sigma^{(i)})$.

A single population

For simplicity, we will start with one population and generalise later. Consider some equations of motions for the dynamics of a single neuron and the corresponding density dynamics

$$\begin{aligned} \dot{x} &= f(x,u) + \Gamma(x) \\ \dot{q} &= -\nabla \cdot fq + \nabla \cdot D\nabla q \end{aligned} \tag{7}$$

Here, we have introduced an exogenous input μ that exerts its effect through the flow (e.g., pre-synaptic input from another population causing a depolarisation and change in voltage). From these equations we can derive the equations of motion for the sufficient statistics (mean and covariance) of the ensemble density, $q(x) = N(\mu, \Sigma)$.

$$\begin{aligned} \dot{\mu}_i &= \int x_i \dot{q}(x)dx \\ &= \int f_i(x) q(x) dx \end{aligned} \tag{8}$$

$$\begin{aligned} \dot{\Sigma}_{ij} &= \int \bar{x}_i \bar{x}_j \dot{q}(x) dx \\ &= \int (\bar{x}_j f_i(x) + \bar{x}_i f_i(x)) q(x) dx + D_{ij} + D_{ij} \end{aligned}$$

Where $\bar{x}_i = (x_i - \mu_i)$ represent perturbations from the mean of the i-th state. These equalities can be verified using integration by parts; for example, with a single state we have

$$\begin{aligned} \dot{\mu} &= \int -x\partial_x(fq)dx + \int x\partial_x D\partial_x q dx \\ &= -[xf(x)q(x)]_{-\infty}^{\infty} + \int f(x)q(x)dx + [xD\partial_x q]_{-\infty}^{\infty} - \int D\partial_x q dx. \\ &= \int f(x)q(x)dx \end{aligned} \tag{9}$$

Here, we have used the fact that $q(x) = \partial_x q(x) = 0 : x \to \pm\infty$ is a proper density. The dynamics of the sufficient statistics in Eq. (8) are intuitively sensible; the rate of change of the mean of any state is the expected flow, in the direction of that state. Similarly, the variance only stops changing when dispersion due to random fluctuations is balanced by contraction due to flow. This contraction is proportional to the negative correlation between flow and the distance from the mean. This perspective can be made explicit by writing Eq. (8) as

$$\begin{aligned} \dot{\mu}_i &= \langle f_i(x) \rangle_q \\ \dot{\Sigma}_{ij} &= \langle \bar{x}_j f_i(x) + \bar{x}_i f_j(x) \rangle_q + D_{ij} + D_{ji} \end{aligned} \tag{10}$$

We can now exploit the fixed-form (Laplace) assumption about the ensemble density by rewriting Eq. (10) in terms of its sufficient statistics, using an expansion of the flow around the expected state

$$f_i(x) = f_i(\mu,u) + \sum_j \frac{\partial f_i}{\partial x_j}\bar{x}_j + \frac{1}{2}\sum_{jk}\frac{\partial^2 f_i}{\partial x_j \partial x_k}\bar{x}_j\bar{x}_k + \ldots \tag{11}$$

Under Gaussian assumptions $\langle \bar{x}_i \rangle_q = 0$ and $\langle \bar{x}_i \bar{x}_j \rangle_q = \Sigma_{ij}$ and we get

$$\begin{aligned} \dot{\mu}_i &= f_i(\mu,u) + \frac{1}{2}\sum_{jk}\frac{\partial^2 f_i}{\partial x_j \partial x_k}\Sigma_{kj} \\ \dot{\Sigma}_{ij} &= \sum_k \frac{\partial f_i}{\partial x_k}\Sigma_{jk} + \sum_k \frac{\partial f_i}{\partial x_k}\Sigma_{ik} + D_{ij} + D_{ji} \end{aligned} \tag{12a}$$

This can be expressed more compactly in matrix form

$$\begin{aligned} \dot{\mu}_i &= f_i(\mu,u) + \frac{1}{2} tr(\Sigma \partial_{xx} f_i) \\ \dot{\Sigma} &= \partial_x f \Sigma + \Sigma \partial_x f^T + D + D^T \end{aligned} \tag{12b}$$

This is a key expression because it allows us to formulate population dynamics, under the Laplace assumption, knowing only the flow, its gradient and curvature (f_i, ∂_x, f_i, ∂_{xx}, f_i) at the expected state. Furthermore, we have circumnavigated the problem of integrating the density at every point in state-space to integrating a small number of sufficient statistics for each population. Eq. (12a,b) is instructive because it shows explicitly how the first and second moments of the density depend on each other; the variance affects the mean when and only when the curvature (second derivative) of the flow is non zero. This will always be the case if the equations of motion are nonlinear in the states. Similarly, the effect of the mean on the variance depends on nonlinear dynamics because the gradients in the second equality above will only change with the mean, when the curvature is non zero.

Interestingly, the form of neuronal dynamics implicit in Eq. (5) is linear in the states; in other words, $\partial_{xx} f_i = 0$. Eq. (12a,b) shows that the dynamics of the mean do not depend on the covariance and a neural-mass model is sufficient to model density dynamics. Below, we will consider a nonlinear conductance-based model where $\partial_{xx} f_i \neq 0$, which means there is a potential role for dispersion. Finally, Eq. (12a,b) shows that if we approximate the ensemble density with a point mass we recover the original equations of motion for a single neuron; i.e., if $\Sigma = 0$ then the dynamics are completely specified in Eq. (12b) by $\dot{\mu}_i = f_i(\mu,u)$. This is a *neural-mass model* and precludes interactions among moments of the population density.

Coupling different populations

Above, we treated each member of the neuronal population as evolving independently of the others, as if we were modelling a 'gas' of neurons. However, real neurons are connected and influence each other. We now consider mean-field equations for a set of m coupled populations that accommodate these influences. Under mean-field coupling each neuron 'senses' the states of all neurons in one or more populations. The ensuing effects can be formulated by making the motion of each neuron a function of population densities and, implicitly, their sufficient statistics, $\mu = \mu^{(1)}, \ldots, \mu^{(m)}$ and $\Sigma = \Sigma^{(1)}, \ldots, \Sigma^{(m)}$

$$\dot{x}^{(i)} = f\left(x^{(i)}, u, \mu, \Sigma\right) + \Gamma(x). \tag{13}$$

This couples the microscopic evolution of each neuron to macroscopic density dynamics within and between populations. These mean-field effects basically change the pattern of flow within a population's state-space. The corresponding density dynamics of the j-th population are now

$$\begin{aligned} \dot{\mu}_i^{(j)} &= f_i^{(j)}(\mu,\Sigma,u) + \frac{1}{2} tr\left(\Sigma^{(j)} \partial_{xx} f_i^{(j)}\right) \\ \dot{\Sigma}^{(j)} &= \partial_x f^{(j)} \Sigma + \Sigma \partial_x f^{(j)T} + D^{(j)} + D^{(j)T} \end{aligned} \tag{14}$$

Notice that the terms involving gradients and curvatures pertain only to the population in question. This is because $\partial f^{(j)} / \partial x^{(i)} = 0 : \forall i \neq j$; in other words the motion in one population $f^{(i)} = f(x^{(i)}, \mu, \Sigma)$ depends only on the density on the states of others, not the states *per se*. Before turning to a specific example we consider the outputs or responses of these systems.

Observed responses

In subsequent work, we will use the density dynamics above as the basis of a dynamic causal model (DCM) of observed data. This requires one to specify how the density maps to observed responses, such as the electroencephalogram (EEG) or blood oxygen level dependent (BOLD) signals in functional magnetic responses imaging (fMRI). Generally, these observations are generated by an average $\eta(\mu,\Sigma) = \langle g(x^{(i)})\rangle_q$; we will use this below. For fMRI $g(x^{(i)}) = H\left(x_j^{(i)} - \tau\right)$ may be a Heaviside or threshold function of depolarisation to reflect synaptic firing. Under the Laplace assumption the expected firing rate $\eta(\mu, \Sigma)$ becomes a sigmoid [error] function of the mean depolarisation. We will pursue the use of density-based DCMs elsewhere (see also Harrison et al., 2005).

Application to a conductance-based model

In this section, we apply the Laplace approximation to a model we have used in previous papers (Friston et al., 2003), which has a complexity intermediate between simple integrate-and-fire models and Hodgkin-Huxley models. This involves specifying the equation of motion and implicitly their gradients and curvatures. These quantities specify the density dynamics in terms of sufficient statistics under the Laplace assumption. Finally, we will look at some special cases that will be compared in the final section of this paper.

The equations of motion

The neuronal dynamics of any given population considered here conform to a simplified Morris and Lecar (1981) model, where the states $x^{(i)} = \left\{V^{(i)}, g_1^{(i)}, g_2^{(i)}, \ldots\right\}$ comprise transmembrane potential and a series of conductances corresponding to different types of ion channel. The dynamics are given by the stochastic differential equations

$$\begin{aligned} C\dot{V}^{(i)} &= \textstyle\sum_k g_k^{(i)}\left(V_k - V^{(i)}\right) + I + \Gamma_V \\ \dot{g}_k^{(i)} &= \kappa_k^{(i)}\left(\varsigma_k^{(i)} - g_k^{(i)}\right) + \Gamma_k \end{aligned} \tag{15}$$

These equations of motion constitute a model for a single neuron and, when solved simultaneously for an ensemble of neurons, furnish an ensemble model. They are effectively the governing equations for a parallel resistance-capacitance circuit; the first says that the rate of change of transmembrane potential (times capacitance, C) is equal to the sum of all currents across the membrane (plus exogenous current, $I = u$). These currents are, by Ohm's law, the product of potential difference between the voltage and reversal potential, V_k for each type of conductance. These currents will either hyperpolarise or depolarise the cell, depending on whether they are mediated by inhibitory or excitatory receptors respectively (i.e., whether V_k is negative or positive). Conductances change dynamically with a characteristic rate constant κ_k and can be regarded as the number of open channels. Channels open in proportion to pre-synaptic input ς_k and close in proportion to the number open. The pre-synaptic input corresponds to the expected firing rate in another population, times a coupling parameter γ_{ij}^k for the k-th conductance

$$\begin{aligned} \varsigma_k^{(i)} &= \textstyle\sum_j \gamma_{ij}^k \int q(V^{(j)}) H(V^{(j)} - V_R) dV^{(j)} \\ &= \textstyle\sum_j \gamma_{ij}^k \sigma\left(\mu_V^{(j)} - V_R, \Sigma^{(j)}\right) \end{aligned} \tag{16}$$

where $H(\cdot)$ is a Heaviside function and the sigmoid function $\sigma(\cdot)$ is a cumulative density on the depolarisation; see Marreiros et al. (2008) and Eq. (17) below. The form of Eq. (16) is motivated in detail in Marreiros et al. (2008; Eq. (6)).

The coupling parameters specify connectivity among populations. Furthermore, they can be used to ensure that each population couples to one and only one conductance type (i.e., each population can only release one sort of neurotransmitter). Generally, one would model a neuronal network of areas, where each area comprises two or more populations. This engenders the distinction between intrinsic and extrinsic connections, which couple populations within and between brain areas. In this paper, we restrict ourselves to a single area and intrinsic connections; however, there is no mathematical distinction between intrinsic and extrinsic connections. The firing in source populations is a Heaviside or threshold function of depolarization where the threshold, V_R determines the proportion of afferent cells firing. Under the mean-field assumption, this input is a function of the population density of the source and under the Laplace assumption this function is simply the Gaussian cumulative density

$$\sigma(\mu,\Sigma) = (2\pi\det(\Sigma))^{\frac{1}{2}} \int_{-\infty}^{\mu} \exp\left(-\tfrac{1}{2}x^T\Sigma^{-1}x\right)dx \tag{17}$$

and is a function of the source's sufficient statistics. These equations constitute $f^{(i)} = f(x^{(i)}, u, \mu, \Sigma)$ of the previous section and are sufficient to elaborate a mean-field approximation under the Laplace assumption using Eq. (14); where (dropping the population superscript for clarity)

$$\begin{aligned} f &= \begin{bmatrix} \frac{1}{C}\sum_k g_k(V_k - V) + \frac{1}{C}I \\ \kappa_1(\varsigma_1 - g_1) \\ \kappa_2(\varsigma_2 - g_2) \\ \vdots \end{bmatrix} \\ \partial_x f &= \begin{bmatrix} -\frac{1}{C}\sum_k g_k & \frac{1}{C}(V_1 - V) & \frac{1}{C}(V_2 - V) & \cdots \\ 0 & -\kappa_1 & 0 & \\ 0 & 0 & -\kappa_2 & \\ \vdots & & & \ddots \end{bmatrix} \\ \partial_{xx} f_V &= \begin{bmatrix} 0 & -\frac{1}{C} & -\frac{1}{C} & \cdots \\ -\frac{1}{C} & 0 & 0 & \\ -\frac{1}{C} & 0 & 0 & \\ \vdots & & & \ddots \end{bmatrix} \\ \partial_{xx} f_g &= 0 \end{aligned} \tag{18}$$

Note that the curvature has a simple form because the equations of motion are second order in only voltage and conductance. An example of the expressions for the ensuing motion of the sufficient statistics $\lambda^{(i)} = \{\mu^{(i)}, \Sigma^{(i)}\}$ from Eq. (14) and the corresponding Jacobian, $\partial_\lambda \cdot \lambda$ are provided in Fig. 1, for two populations. This figure provides an iconic summary of how different quantities affect each other. For example, the variances affect only the mean depolarisation, in inverse proportion to the capacitance. These equations are not necessary to integrate the sufficient statistics; we only derived the subset of equations shown in the figure for didactic purposes. In practice, these derivatives are evaluated numerically, given the user specified equations of motion.

Some special cases

Before assessing the accuracy of the Laplace scheme we will consider some special cases of Eq. (14). The first obtains if we assume $\Sigma^{(i)}$ is fixed for all populations. Because the covariance is fixed, we only have to integrate the ensemble mean; furthermore because the curvature is constant (voltage) or zero (conductance), this entails an extra decay term for voltage, giving density dynamics of the form

$$\begin{aligned} \dot{\mu}_V^{(i)} &= f_V^{(i)}(\mu,\Sigma,u) + \tfrac{1}{2}tr\left(\Sigma^{(i)}\partial_{xx}f_V^{(i)}\right) \\ \dot{\mu}_g^{(i)} &= f_g^{(i)}(\mu,\Sigma,u) \end{aligned} \tag{19}$$

This corresponds to a neural-mass model with decay and will be used for comparative analyse in the next section. Finally if we

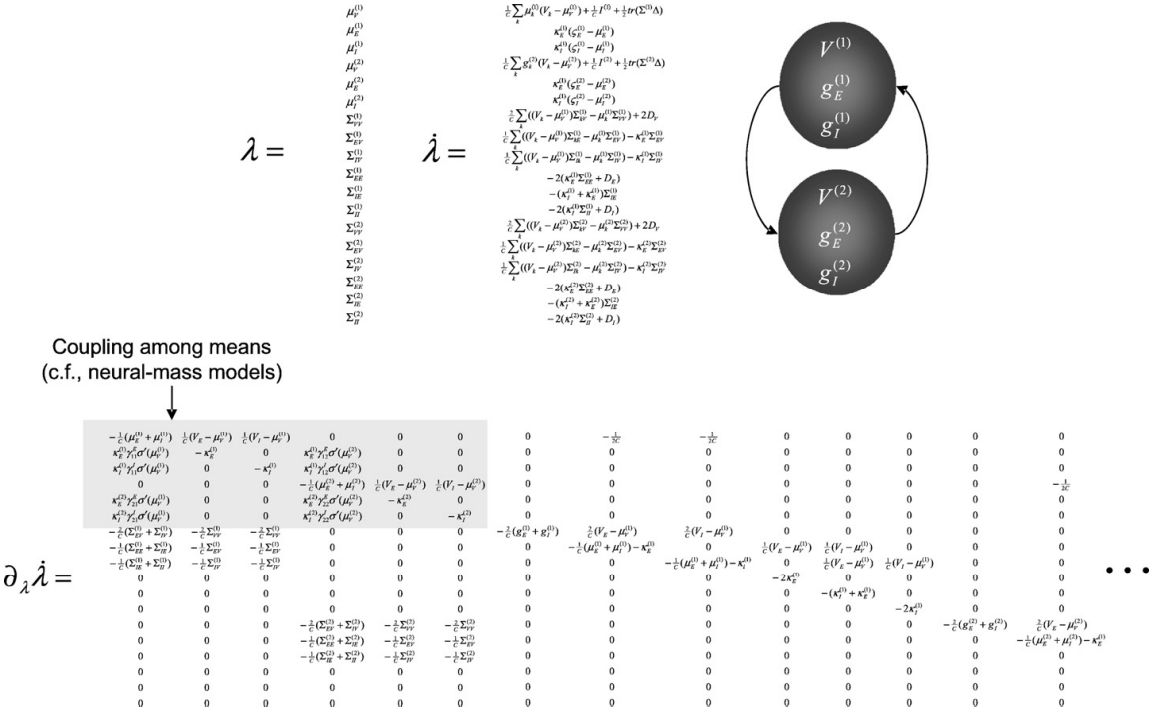

Fig. 1. Expressions for the motion of the sufficient statistics $\lambda^{(i)}=\mu^{(i)}$, $\Sigma^{(i)}$ (mean and variance) and the corresponding Jacobian for two populations that conform to simplified Morris-Lecar-like dynamics. The grey area in the Jacobian covers terms that link mean states to each other and are this considered in neural-mass reductions of full mean-field models.

further assume that $\Sigma^{(i)}$ is spherical (i.e., all off-diagonal terms are zero) then this decay terms disappears because the leading diagonal of $\partial_{xx}f_V^{(i)}$ is zero. In this instance, the dynamics reduce to the original equations of motion because we can ignore the second-order statistics completely

$$\dot{\mu}_k^{(i)} = f_k^{(i)}(\mu,\Sigma,u). \tag{20}$$

This is a conventional neural-mass model with the usual sigmoid activation function. This function depends on the variance (see Eq. (15)), which we assume is fixed. Note that this provides another perspective on the parameterisation of the sigmoid activation function in classical neural-mass models (c.f. Eq. (5) and the derivations in Marreiros et al., 2008). In the next section we will compare the Laplace (Eq. (14)) and neural-mass (Eq. (19)) approximations in terms of modelling evoked neuronal transients.

Summary

We are now in a position to compare and contrast ensemble models of neuronal populations with mean-field (MFM) and neural-mass (NMM) approximations. Ensemble models (Eq. (15)) provide the trajectories of many neurons to form a sample density of population dynamics. The MFM is obtained by a mean-field and a Laplace approximation to these densities (Eq. (14)). The NMM is a special case of the mean-field model in which we ignore all but the first moment of the density (i.e., the mean or mode). In other words, the NMM discounts dynamics of second-order statistics (i.e., variance) of the neuronal states. The mean-field models allow us to model interactions between the mean of neuronal states (e.g., firing rates) and their dispersion or variance over each neuronal population modelled (c.f., Harrison et al., 2005). The key behaviour we are interested in is the coupling between the mean and variance of the ensemble, which is

Table 1
Overview of the three models: Ensemble, Mean-Field model (MFM) and Neural-Mass model (NMM)

Model	Description	Equation
Ensemble	Stochastic differential equation that describes how the states evolve as functions of each other and some random fluctuations	$dx = f(x,u)dt + \sigma dw$ (Eq. (2))
MFM	Differential equation that describes how the density evolves as functions of mean and covariance. Resulting from a mean-field and Laplace approximations of the ensemble model	$\dot{\mu}_i^{(j)} = f_i^{(j)}(\mu,\Sigma,u) + \frac{1}{2}tr\left(\Sigma^{(j)}\partial_{xx}f_i^{(j)}\right)\cdot\Sigma^{(j)} = \partial_x f^{(j)}\Sigma + \Sigma\partial_x f^{(j)T} + D^{(j)} + D^{(j)T}$ (Eq.14)
NMM	Differential equation that describes how the density evolves as a function of the mean. Obtained by fixing the covariance of the MFM	$\dot{\mu}_i^{(j)} = f_i^{(j)}(u,\Sigma,\mu) + \frac{1}{2}tr\left(\Sigma^{(j)}\partial_{xx}f_i^{(j)}\right)\cdot\Sigma^{(j)} = 0$ (Eq. (19))

For a detailed description of the equations see main text.

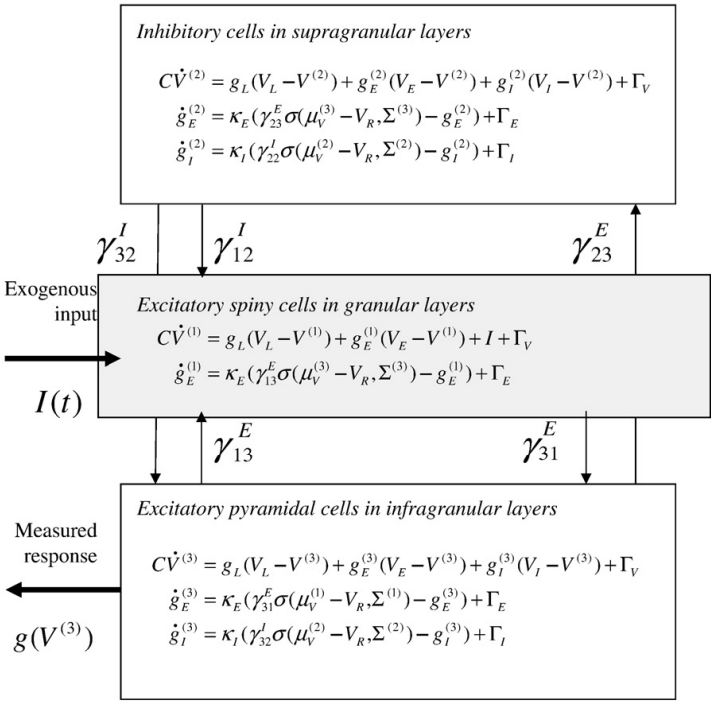

Fig. 2. Source model with layered architecture comprising of three neuronal populations (Spiny-stellate, Interneurons, and Pyramidal cells), each of which has three different states (Voltage, Excitatory, and Inhibitory conductances).

lost in the NMM. The different models and their mathematical representations are summarised in Table 1.

Neural-mass vs. mean-field models

In this section, we examine the accuracy of the Laplace approximation to density dynamics, in relation to the true density dynamics that obtain by integrating the trajectories of a real but finite-sized population. We will also take the opportunity to highlight the difference between the Laplace approximation and neural-mass simplifications. In what follows, we examine the response of three populations connected to emulate the source model for electromagnetic responses we use in DCM for ERPs (David et al., 2006, Kiebel et al., 2006). Each electromagnetic source comprises two excitatory populations and an inhibitory population. These are taken to represent input cells (spiny stellate cells in the granular layer of cortex), inhibitory interneurons (allocated somewhat arbitrarily to the superficial layers) and output cells (pyramidal cells in the deep layers). The deployment and intrinsic connections among these populations are shown in Fig. 2 and the parameters are provided in Table 2.

In this model, we use three conductance types: leaky, excitatory and inhibitory conductance. This gives, for each population

$$\begin{aligned} C\dot{V}^{(i)} &= g_L(V_L - V^{(i)}) + g_E^{(i)}(V_L - V^{(i)}) + g_E^{(i)}(V_I - V^{(i)}) + I + \Gamma_V \\ \dot{g}_E^{(i)} &= \kappa_E\left(\varsigma_E^{(i)} - g_E^{(i)}\right) + \Gamma_E \\ \dot{g}_I^{(i)} &= \kappa_I\left(\varsigma_I^{(i)} - g_I^{(i)}\right) + \Gamma_I \\ \varsigma_k^{(i)} &= \textstyle\sum_j \gamma_{ij}^k \sigma\left(\mu_V^{(j)} - V_R, \Sigma^{(j)}\right) \end{aligned} \tag{21}$$

Notice that the leaky conductance does not change, which means the states reduce to, $x^{(i)} = \left\{V^{(i)}, g_E^{(i)}, g_I^{(i)}\right\}$. Furthermore, for simplicity, we have assumed that the rate-constants, like the reversal potentials are the same for each population. The excitatory and inhibitory nature of each population is defined entirely by the specification of the non-zero intrinsic connections γ_{ij}^k (see Fig. 2). The resulting sparse connectivity means that not all populations have all conductances.

Simulations

In what follows, we examine the response of this three-population source to exogenous input using the Laplace and neural-mass approximations. We first compare the analytic approximations based on the mean-field (Eq. (14)), with the sample density of responses from simulated neuronal ensemble (Eq. (15)). We present more comprehensive characterisations, comparing predicted responses under mean-field and neural-mass models to transient and sustained input. Our aim was to (i) evaluate the Laplace approximation in relation to the response obtained by integrating the original stochastic equation of motions and (ii) to compare the Laplace approximation (Eq. (14)) with the neural-mass model (Eq.

Table 2
Parameter values for all models used in this paper

Parameter	Physiological interpretation	Value
g_L	Leaky conductance	1 mV
$\tau_{E,I} = 1/K_{E,I}$	Postsynaptic rate constants	4 ms, 16 ms
$\gamma_{31}^E, \gamma_{13}^E, \gamma_{23}^E, \gamma_{32}^I, \gamma_{32}^I$	Intrinsic connectivity	1, 0.5, 1, 0.5, 2
V_L, V_E, V_I	Reversal potential	-70 mV, 60 mV, -90 mV
V_R	Threshold potential	-40 mV

(19)) to assess the need for population covariance as part of the model.

Ensemble dynamics

In the first simulations, we examined population responses to an impulse or burst of afferent input. This can be regarded as a simple evoked response. We integrated the equations of motion (Eq. (15)) for the three population model of Fig. 2, with 64 neurons per population. To integrate the stochastic differential equations, we added a random normal variate to the states of each neuron, at each time step Δt sampled from a Gaussian density with variance, $2D\Delta t$. The ensuing impulse responses are shown in Fig. 3, in terms of the depolarisation of pyramidal cells. Because we used a relatively small ensemble of neurons there are some (but not marked) finite-size effects: Finite-size effects are seen when approximating the response of a large ensemble with the response of a small number of neurons (see Mattia and Del Giudice, 2004; Doiron et al., 2006 and Galán et al., 2007 for a discussion of finite element methods in characterising the behaviour of neuronal ensembles). Critically, the random fluctuations due to the Wiener processes lead to different trajectories (Fig. 3; middle panel), which provide a sample density for the population dynamics. This can be summarised in terms of its mean and a 90% confidence interval, over peristimulus time (Fig. 3, lower panel). The key thing to observe here is that the dispersion is not stationary; it changes with time. Specifically, when the states are changing quickly around the peak response, the dispersion of states is much smaller than when the ensemble is at baseline. It is this change in dispersion that is discounted by conventional NMMs.

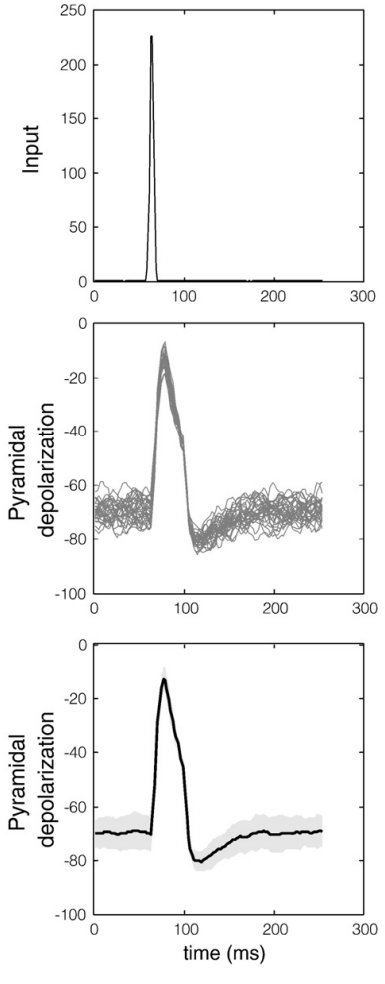

Fig. 3. Top: Exogenous input. Middle: Integrated response (64 neurons) of the pyramidal population, where the spike is driving the neuronal source through intrinsic connections (Fig. 2). Bottom: Summary of the density over trajectories in terms of their mean (solid line) and a 90% confidence interval (grey region).

Fig. 4 shows the integrated responses of this ensemble of neurons, for all states and populations. The red arrows show the main causal influences that couple different populations. These are the mean-field effects of depolarisation in one population increasing the excitatory or inhibitory conductance of another (through intrinsic connections). This, in turn leads to depolarisation or hyperpolarisation of the target population. The configuration of intrinsic connections means that input, which enters at the spiny stellate population, may only be expressed ten or more milliseconds later in other populations. It is these slow population effects we want to approximate.

We solved the population dynamics to give the MFM and the NMM approximations to the impulse responses in Fig. 4. For the MFM, the mean and dispersion of the state dynamics were computed by solving Eq. (14) for the same model and input used above. The NMM dynamics were obtained by fixing the dispersion of the MFM to its steady-state value (in the absence of input); this is the stationary solution to Eq. (14). In Fig. 5, we plot the responses from all three models. One can see that the mean of the trajectories are similar for all models. Although one can see that after the peak, the mean response of the MFM response is more like the ensemble model than the NMM response. Furthermore, like the ensemble model, the MFM dispersion changes over time, while the dispersion of the NMM is constant. For more complex source models these small differences may have significant repercussions, if the dynamics of the mean depend on dispersion. We will see an example of this later. The MFM appears to overestimate the dispersion in comparison to the ensemble model; however, this is probably due to finite size effects.

Comparing MFM and NMM predictions

Using the above model, we compared the MFM and the NMM responses using exogenous inputs that varied in amplitude and were transient or sustained. The results of these simulations are shown in Fig. 6, in terms of pyramidal cell population depolarisation. With transient inputs we found that both the MFM and the NMM predicted a similar response. The two models respond with a short-lived burst of activity that increased with input amplitude and showed a plateau around 60 μA. When the input exceeds 90 μA, the response under both models become biphasic, with a second peak that lasted for about 30 ms. With sustained input, both models show complex nonlinear behaviour for input amplitudes greater than 24 μA. However, for input amplitude values greater than 50 μA the response patterns of the two models are very different. The MFM shows a sustained oscillatory or limit-cycle behaviour that is largely unaffected by further increases in input. In contrast, the NMM returns to a fixed level of depolarization (a fixed-point attractor) after about 200 ms; this illustrates that the MFM retains key nonlinearities and can exhibit bifurcations that are structurally distinct from the NMM. In short, one observes subtle but potentially important differences between the two models, which may have important implications for generative models of observed neuronal responses.

A quantitative characterisation

To quantify neuronal responses to sustained input under the MFM, we used frequency analyses and mean spiking responses. We focussed

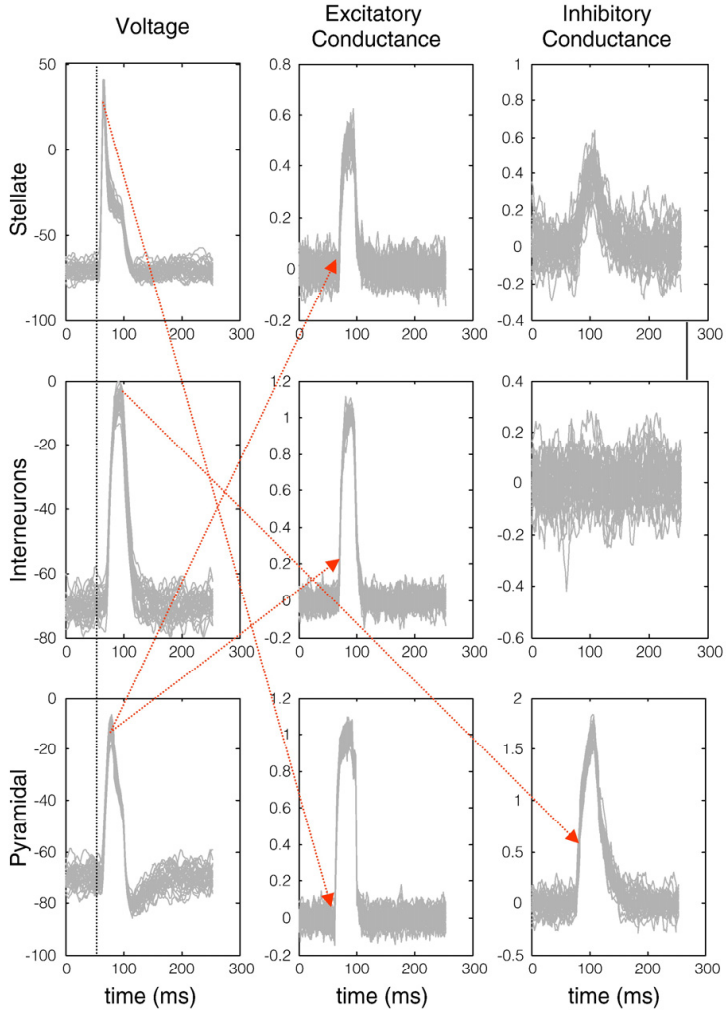

Fig. 4. Ensemble model responses for the three neuronal populations (stellate, interneurons, pyramidal) over their three different states (voltage, excitatory and inhibitory conductance). The red lines correspond to the causal influences mediated by intrinsic connections that convey means-field effects (from voltage to conductances). The vertical broken line is aligned to the exogenous input that arrives at 64 ms.

on the pyramidal population, which represents the principal (output) cells in cortex and are the predominant source of electromagnetic signals that are observed empirically. The results of these analyses are shown in Fig. 7 using the same model and range of sustained input as above. It can be seen that the spectral responses are greatest between about 8 and 16 Hz, for input amplitudes between 25 and 45 μA (Fig. 7A). In this range, the peak frequency increases almost linearly with amplitude. In Fig. 7B we look in more detail at the MFM spectral response profile at input amplitudes of 32 and 64 μA. These two input levels fall into two different regimes of the spectral response (broken lines in Fig. 7A). For the 32 μA input there is a pronounced alpha peak at ~10 Hz, for the 64 μA input, the spectrum has a small beta peak around 24 Hz.

We next looked at how the population firing response probability scales with input amplitude. Fig. 7C shows that a response emerges, after about 100 ms, at about 25 μA input amplitude and shows nonlinear behaviour over time; for higher input amplitudes, the activity oscillates at a constant frequency. This response pattern is very similar to the depolarization (Fig. 6), because firing rate is a nonlinear function of the density on pyramidal depolarization. The ensuing input-firing rate curve (averaged over peristimulus time) shows a highly nonlinear behaviour, with no firing below a threshold of 20 μA and progressive increases until the firing saturates at input amplitudes of about 50 μA (Fig. 7D).

This sort of simulation demonstrates that the limit-cycle attractor of the MFM can be exploited to study the relationship between oscillatory dynamics and mean levels of firing. In this instance, the model suggests that high firing rates, induced by sustained inputs, will be expressed in the context of higher frequencies in a desynchronised or 'activated' EEG. This is entirely consistent with empirical observations (e.g., Kilner et al 2005 and references therein). More generally, this simple simulation shows that the nature of responses predicted by mean-field and neural-mass models of exactly the same neuronal system can differ profoundly in terms of

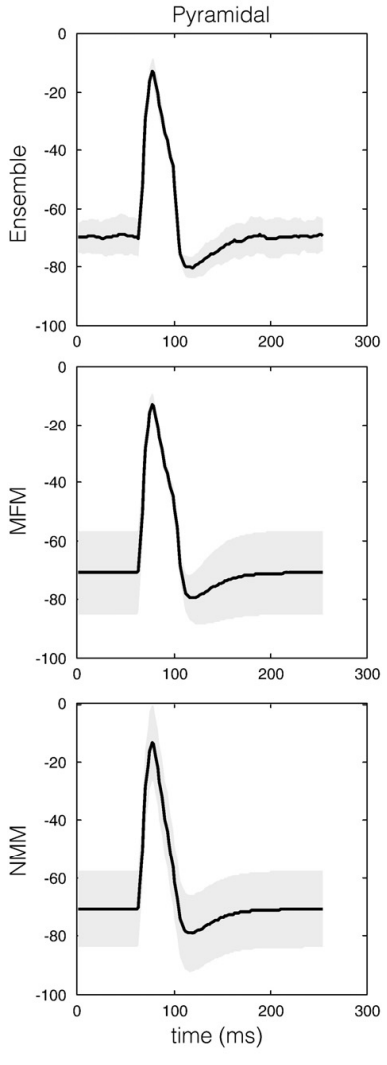

Fig. 5. Population response of the pyramidal cells for the three models: ensemble model, mean-field model and neural-mass model. One can see differences for the mean (solid lines) and the dispersion (grey regions) of the trajectories.

the dynamics they support. Here, the addition of extra variables encoding population covariance leads to oscillations, under sustained input that are not predicted by a reduced neural-mass model. In principle, this means that mean-field DCMs of evoked and induced responses may provide better models of empirical data. We pursue the theme of nonlinearity and limit-cycles in the final simulations, which look at nested oscillations.

Modelling nested oscillations and phase-synchronisation

Nonlinear coupling between distinct brain regions are observed, predominantly as interactions between low and high frequencies. These nonlinear influences are thought to mediate top-down modulation, 'attentional' and other context-defining functions (Kopell et al., 2000; Von Stein et al., 2000; Varela et al., 2001; Canolty et al., 2006). Two principal forms of cross-frequency phase interactions have been recognized: '*n:m* phase synchrony', which indicates amplitude-independent phase-locking of *n* cycles of one oscillation to *m* cycles of another oscillation (Tass et al., 1998; Palva et al., 2005); and 'nested oscillations', which reflect the locking of the amplitude fluctuations of faster oscillations to the phase of a slower oscillation (Vanhatalo et al., 2004; Canolty et al., 2006, Penny et al., 2008). Nested oscillations have been observed in both the human brain and rat hippocampus (Chrobak and Buzsaki, 1998; Mormann et al., 2005); they have been proposed to underlie the discrete nature of perception and the capacity of working memory (Penny et al., 2008), as well as playing a role in sleep (Steriade, 2006) and olfaction (Kepecs et al., 2006). There many studies which rest on cross-frequency coupling, for example (Lisman and Idiart, 1995; Fukai, 1999; Hocking and Levy, 2007; Haenschel et al., 2007).

Motivated by these findings, we reproduced nested oscillations using our three-population source (Fig. 2). We drove the neuronal source with a slow sinusoidal input to elicit periods of bursting in the inhibitory population. This produced phase-amplitude coupling, most notably between the inhibitory population and the spiny population that was driven by the low-frequency input. The bursting and concomitant nested oscillations are caused by nonlinear interactions between voltage and conductance, which are augmented by coupling between their respective means and dispersions. Fig. 8 shows the predicted responses from the MFM and NMM models. The population responses of the MFM and NMM show clear differences in the number and amplitude of the oscillations per cycle of the low frequency input. Again this illustrates the potential importance of using a MFM (as opposed to a NMM) when modelling nonlinear or quasi-periodic dynamics, like nested oscillations. This simulation is another illustration of how small differences between models can have large effects on the nature of predicted neuronal responses.

Discussion

The purpose of this work was to describe a generic approach to modelling dynamics in neuronal populations. Our work is motivated by the observation that neural-mass approaches, currently used as generative models for observed data (David et al., 2006), are a limiting case of mean-field models. In other words, they consider only the first moment of the density for each population, which is a special case of the more general ensemble density formulation. In this paper we augmented the neural-mass model with quantities that encode population dispersion to furnish mean-field models that capture full density dynamics.

The high dimensionality and complexity of Fokker-Planck formalisms can be reduced with a mean-field approximation, which describes the evolution of separate ensembles coupled by mean-field effects. By parameterising the densities in terms of their sufficient statistics, the partial differential equations can be reduced to ordinary differential equations describing the evolution of its sufficient statistics or moments (Table 2). In this way, we obtained a key equation (Eq. (14)), which formulates population dynamics, using only the flow, its gradient and curvature, at the mean state. This expression shows explicitly how the first and second moments of the density depend on each other; the variance affects the mean if and only if the curvature (second derivative) of the flow is non zero. This will be the case if the equations of motion are nonlinear in the neuronal states. Similarly, the effect of the mean on the variance depends on nonlinear dynamics because the gradients will only change with the mean, when the curvature is non zero.

We have looked at the neuronal response of a particular but ubiquitous model (Fig. 2) in terms of the mean and the dispersion of its underlying neuronal states (Figs. 3 and 4). We established the validity of the Laplace approximation by comparing the response of a simulated ensemble of neurons to the response under the Laplace and neural-mass assumptions. The key behaviour we were interested in

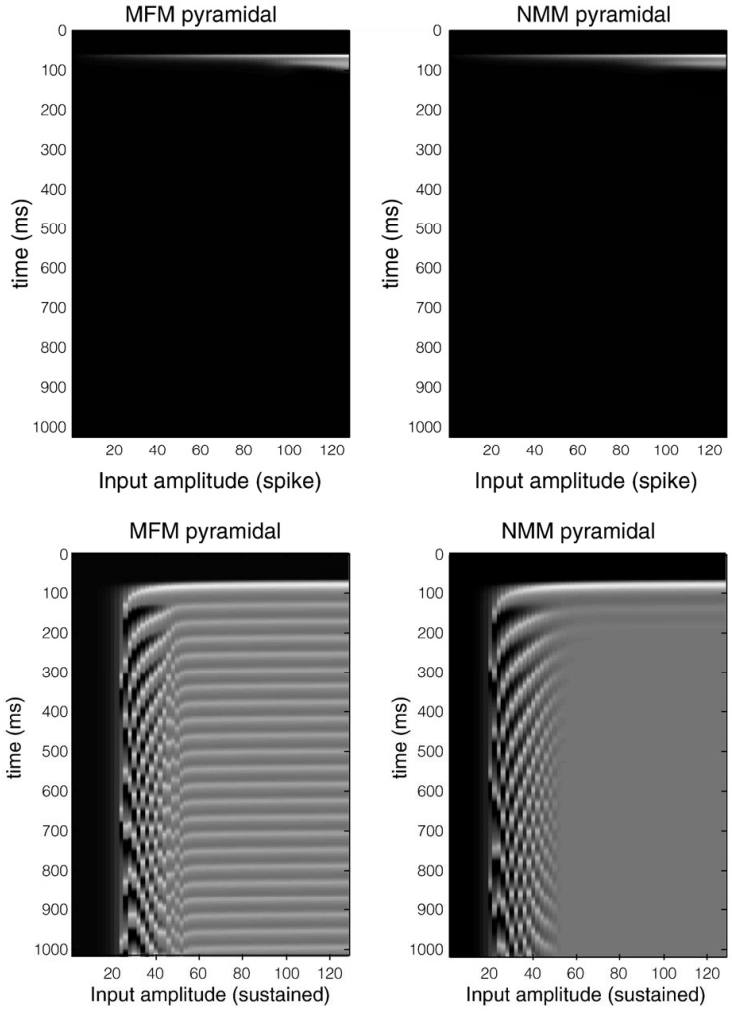

Fig. 6. (Left): Pyramidal population response (depolarization) under the mean-field model to varying levels of input. (Right): Equivalent pyramidal population response under the neural-mass model. (Top row) transient input at 64 ms; (Lower row) sustained input. The key thing to note is the difference between the predictions of the two models in the lower panels, which show the mean-field model prediction to oscillate at high levels of input. White indicates -10 mV and black -80 mV.

was the coupling between the mean and variance of the ensemble, which is lost in the neural-mass approximations. This enabled us to compare equivalent mean-field and neural-mass models of the same populations and evaluate, quantitatively, the contribution of population variance to shaping population dynamics. The simulations for the Laplace mean-field model, which considers second-order statistics, support a more realistic and plausible model than the neural-mass model. The MFM shows, for an impulse response function, a dynamical behaviour that is more similar than the NMM to the response obtained by integrating the stochastic ensemble dynamics (Fig. 5). Although the NMM enjoys much attention because of its simplicity, it only considers the mean neuronal state and does not consider higher statistics like the variance. We speculate that this simplifying assumption may have implications when trying to invert generative models of real data.

We compared the Laplace approximation (Eq. (14)) with the neural-mass model (Eq. (19)) to assess the role of the population covariance. NMMs, despite their relative simplicity, exhibited complex dynamical behaviour reminiscent of real neuronal responses. However, qualitative differences between MFM and NMM predictions were easy to demonstrate. In particular, we saw that the MFM showed a bifurcation from fixed-point to a limit-cycle attractor, as sustained input levels were increased (Fig. 6). We also looked at the spectral response of the pyramidal population of the mean-field model (Fig. 7). This analysis disclosed the presence of physiologically plausible oscillatory signals in the alpha and beta band and how their relative power changed with activation. Additionally, we presented an interesting example of the quantitative difference between MFM and NMM by reproducing nested oscillation behaviour (Fig. 8). In short, the MFM appeared to represent richer and more complex dynamics. This approach may have potential applications in dynamic causal modelling of imaging studies (M/EEG, fMRI) where one tries to explain the coordinated activity of a large number of neurons.

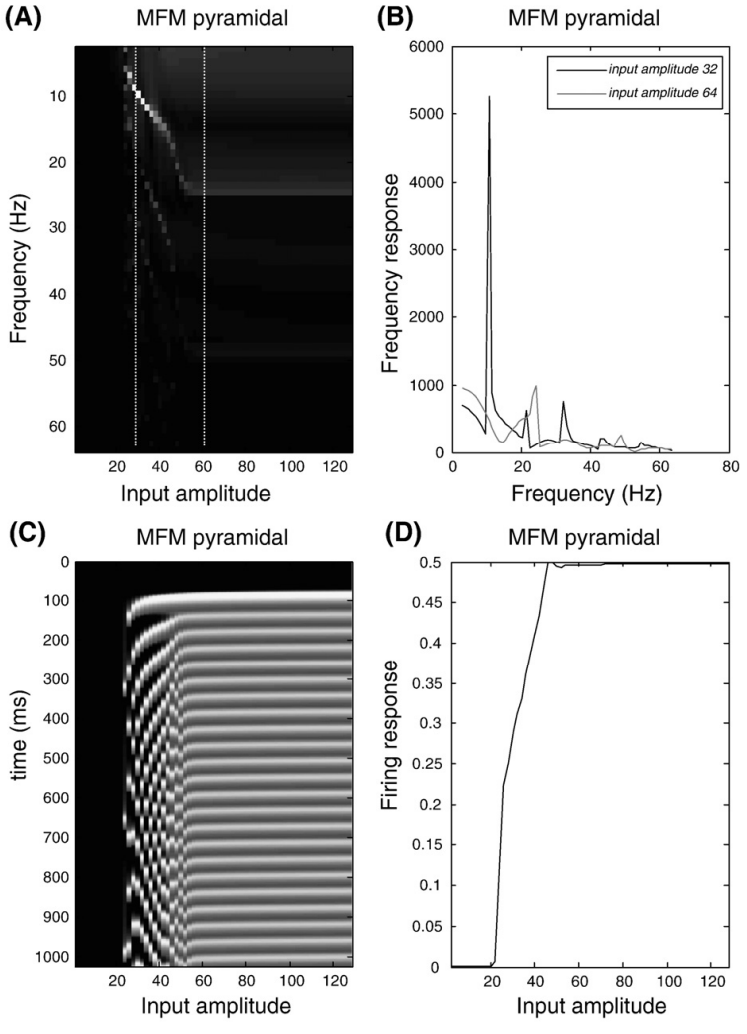

Fig. 7. Mean-field model frequency response for the pyramidal population. (A) Spectral density of response as a function of input amplitude; (B) Spectral density of response for input amplitude of 32 and 64 μA (broken lines in A); (C) Pyramidal firing rates as a function of time and input amplitude; (D) Mean population firing over time as a function of input amplitude.

The Laplace assumption is a common device in statistical physics and finesses the problem of integrating a complicated density by assuming a Gaussian form. In machine learning, it allows one to focus on its sufficient statistics, namely the mean and covariance. In the present context, it allows one to summarise density dynamics using the method of moments (MM; Rodriguez and Tuckwell 1996, 1998). This entails replacing the system of stochastic differential equations with a system of deterministic equations (ODE) representing the dynamics of the means, variances, and covariance of the state variables, i.e., the first and second-order moments of the population density. This is precisely what we have done; namely, derive the ODE for the sufficient statistics of a Gaussian population density, given any set of Fokker-Planck equations that are coupled by phase-functions specifying mean-field effects or effective connectivity.

In related work, Hasegawa has proposed a semi-analytical mean-field approximation, in which the equations of motion for moments were derived for FN and HH ensembles (Hasegawa, 2003a,b). Later he proposed an augmented moment method (AMM; Hasegawa 2004), which relaxes the Gaussian or Laplace approximation (Hasegawa, 2006, 2007). In Deco and Martí (2007), the MM was extended to cover bimodal densities on the state variables; such that a reduced system of deterministic ODEs could be derived to characterise regimes of multistability. The ODEs in Fig. 1 pertain to Morris-Lecar-like neurons and will form the basis of dynamical causal models of empirical EEG and LFP data.

Conclusion

In conclusion, we have derived a generic mean-field treatment of neuronal dynamics, which is based on a Laplace approximation to the ensemble density and is formulated in terms of equations of motion for the sufficient statistics of the ensemble density. We saw how this

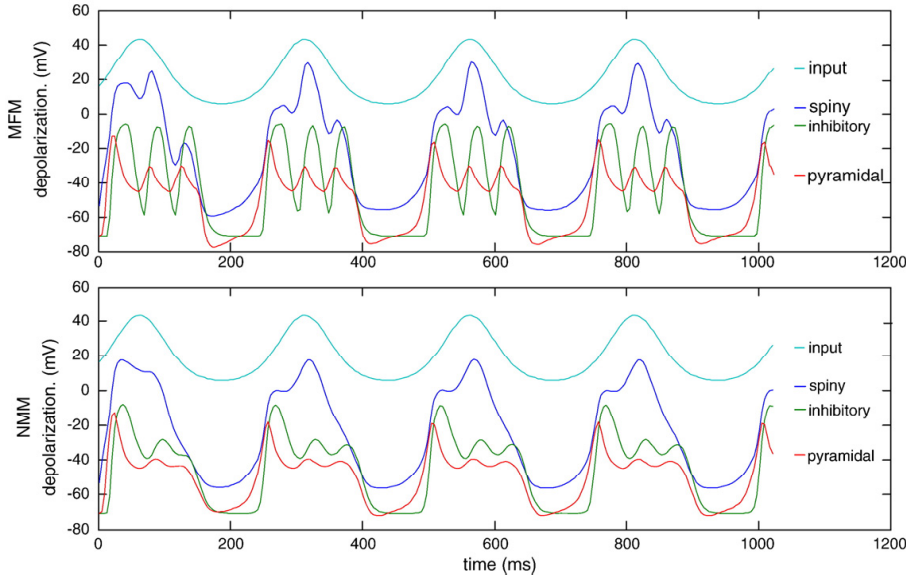

Fig. 8. Nested oscillations in the three-population source driven by slow sinusoidal input for both MFM and NMM. Input is shown in light blue, spiny interneuron depolarization in dark blue, inhibitory interneurons in green and pyramidal depolarization in red. The nonlinear interactions between voltage and conductance produces phase-amplitude coupling in the ensuing dynamics. The MFM shows deeper oscillatory responses during the nested oscillations.

approach reduces to a neural-mass model when the second-order statistics (i.e., variance) of neuronal states is ignored. In a future paper, we will use the Laplace and neural-mass approximations presented here as generative models of electrophysiological responses to sensory input. This paper will use Bayesian model comparison to compare both models and establish whether empirical responses contain evidence for a role of the variance in shaping population dynamics. This framework will allow one to adjudicate between models that include the high-order statistics of neuronal states in predicting EEG time series and may also be important in the context of EEG-fMRI fusion; where power (second-order statistics) in electrical dynamics may be an important predictor of BOLD signals.

Acknowledgments

The Portuguese Foundation for Science and Technology and the Wellcome Trust supported this work.

References

Beurle, R.L., 1956. Properties of a mass of cells capable of regenerating pulses. Philos. Trans. R. Soc. Lond. B. 240, 55–94.
Breakspear, M., Roberts, J.A., Terry, J.R., Rodrigues, S., Mahant, N., Robinson, P.A., 2006. A unifying explanation of primary generalized seizures through nonlinear brain modeling and bifurcation analysis. Cereb. Cortex 16 (9), 1296–1313.
Canolty, R.T., Edwards, E., Dalal, S.S., Soltani, M., Nagarajan, S.S., Kirsch, H.E., Berger, M.S., Barbaro, N.M., Knight, R.T., 2006. High gamma power is phase-locked to theta oscillations in human neocortex. Science 313, 1626–1628.
Casti, A.R., Omurtag, A., Sornborger, A., Kaplan, E., Knight, B., Victor, J., Sirovich, L., 2002. A population study of integrate-and-fire-or-burst neurons. Neural Comput. 14, 957–986.
Chizhov, A., Graham, L., 2007. Population model of hippocampal pyramidal neurons, linking a refractory density approach to conductance-based neurons. Phys. Rev. E 75, 011924.
Chizhov, A.V., Graham, L.J., Turbin, A.A., 2006. Simulation of neural population dynamics with a refractory density approach and a conductance-based threshold neuron model. Neurocomputing 70, 252–262.
Chrobak, J.J., Buzsaki, G., 1998. Gamma oscillations in the entorhinal cortex of the freely behaving rat. J. Neurosci. 18, 388–398.
Chumbley, J.R., Friston, K.J., Fearn, T., Kiebel, S.J., 2007. A Metropolis-Hastings algorithm for dynamic causal models. NeuroImage 38 (3), 478–487.
David, O., Friston, K.J., 2003. A neural-mass model for MEG/EEG: Coupling and neuronal dynamics. NeuroImage 20, 1743–1755.
David, O., Kiebel, S.J., Harrison, L.M., Mattout, J., Kilner, J.M., Friston, K.J., 2006. Dynamic causal modeling of evoked responses in EEG and MEG. NeuroImage 30 (4), 1255–1272.
Dayan, P., Abbott, L., 2001. Theoretical neuroscience. MIT Press, Cambridge, MA.
Deco, G., Martí, D., 2007. Extended method of moments for deterministic analysis of stochastic multistable neurodynamical systems. Physical Review E 75, 031913.
Deco, G., Jirsa, V.K., Robinson, P.A., Breakspear, M., Friston, K.J., 2008. The dynamic brain: from spiking neurons to neural-masses and cortical fields. PLoS-CB 4 (8), e1000092.
Doiron, B., Rinzel, J., Reyes, A.D., 2006. Stochastic synchronization in finite size spiking networks. Phys. Rev. E 74 (rapid comm.).
Eggert, J., van Hemmen, J.L., 2001. Modeling neuronal assemblies: theory and implementation. Neural Comput. 13 (9), 1923–1974.
Elbert, T., Ray, W.J., Kowalik, Z.J., Skinner, J.E., Graf, K.E., Birbaumer, N., 1994. Chaos and Physiology: Deterministic Chaos in Excitable Cell Assemblies. Physiol. Rev. 74 No 1.
Frank, T.D., 2004. Nonlinear Fokker–Planck Equations: Fundamentals and Applications. Springer Series in. Synergetics. Springer, Berlin.
Frank, T.D., Daffertshofer, A., Beek, P.J., 2001. Multivariate Ornstein–Uhlenbeck processes with mean-field dependent coefficients: application to postural sway. Phys. Rev., E Stat. Nonlinear Soft Matter Phys. 63, 011 905.
Freeman, W.J., 1975. Mass Action in the Nervous System. Academic Press, New York.
Friston, K.J., Harrison, L.M., Penny, W.D., 2003. Dynamic causal modelling. NeuroImage 19 (4), 1273–1302.
Friston, K., Mattout, J., Trujillo-Barreto, N., Ashburner, J., Penny, W., 2007. Variational free energy and the Laplace approximation. NeuroImage 34 (1), 220–234.
Fukai, T., 1999. Sequence generation in arbitrary temporal patterns from theta-nested gamma oscillations: a model of the basal ganglia-thalamo-cortical loops. Neural Netw. 12 (7–8), 975–987.
Galán, R.F., Ermentrout, G.B., Urban, N.N., 2007. Stochastic dynamics of uncoupled neural oscillators: Fokker-Planck studies with the finite element method. Phys. Rev. E 76, 056110.
Gerstner, W., Kistler, W.M., 2002. Spiking Neuron Models Single Neurons, Populations. Plasticity Cambridge University Press.
Griffith, J.S., 1963. A field theory of neural nets: I: derivation of field equations. Bull. Math. Biophys. 25, 111–120.
Griffith, J.S., 1965. A field theory of neural nets: II: properties of field equations. Bull. Math. Biophys. 27, 187–195.
Haenschel, C., Ulhaas, P.J., Singer, W., 2007. Synchronous oscillatory and working memory in schizophrenia. Pharmacopsychiatry 40, S54–S61.
Harrison, L.M., David, O., Friston, K.J., 2005. Stochastic models of neuronal dynamics. Philos. Trans. R. Soc. Lond. B Biol. Sci. 360 (1457), 1075–1091.
Hasegawa, H., 2003a. Dynamical mean-field theory of spiking neuron ensembles: Response to a single spike with independent noises. Phys. Rev. E 67, 041903.

Hasegawa, H., 2003b. Dynamical mean-field theory of noisy spiking neuron ensemble: Application to the Hodgkin-Huxley model. Phys. Rev. E 68, 041909.
Hasegawa, H., 2004. Dynamical mean-field approximation to small-world networks of spiking neurons: From local to global and/or from regular to random couplings. Phys. Rev. E 70, 066107.
Hasegawa, H., 2006. N-dependent multiplicative-noise contributions in finite n-unit Langevin models: augmented moment approach. J. Phys. Soc. Jpn. 75, 033001.
Hasegawa, H., 2007. Stationary and dynamical properties of finite N-unit Langevin models subjected to multiplicative noises. Physica, A 374 (2), 585–599.
Haskell, E., Nykamp, D.Q., Tranchina, D., 2001. Population density methods for large-scale modelling of neuronal networks with realistic synaptic kinetics: cutting the dimension down to size. Network 12, 141–174.
Hocking, A.B., Levy, W.B., 2007. Theta-modulated input reduces intrinsic gamma oscillation in a hippocampal model. Neurocomputing 70 (10–12), 2074–2078.
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, B.H., Balaji Kavaipatti, A., Markusson, O., 2001. Evoked potential enhancement using a neurophysiologically-based model. Methods Inf. Med. 40, 338–345.
Jirsa, V.K., Haken, H., 1996. Field theory of electromagnetic brain activity. Phys. Rev. Lett. 77, 960–963.
Kepecs, A., Uchida, N., Mainen, Z., 2006. The sniff as a unit of olfactory processing. Chem. Senses 31, 167–179.
Kiebel, S.J., David, O., Friston, K.J., 2006. Dynamic causal modelling of evoked responses in EEG/MEG with lead-field parameterization. Neuroimage 30, 1273–1284.
Kilner, J.M., Mattout, J., Henson, R., Friston, K.J., 2005. Hemodynamic correlates of EEG: a heuristic. Neuroimage 28 (1), 280–286 Oct 15.
Kincses, W.E., Braun, C., Kaiser, S., Elbert, T., 1999. Modeling extended sources of event-related potentials using anatomical and physiological constraints. Hum. Brain Mapp. 8 (4), 182–193.
Knight, B.W., 2000. Dynamics of encoding in neuron populations: some general mathematical features. Neural Comput. 12, 473–518.
Kopell, N., Ermentrout, G.B., Whittington, M.A., Traub, R.D., 2000. Gamma rhythms and beta rhythms have different synchronization properties. Proc. Natl. Acad. Sci. U. S. A. 97 (4), 1867–1872.
Lopes da Silva, F.H., Hoeks, A., Smits, H., Zetterberg, L.H., 1974. Model of brain rhythmic activity: the alpha-rhythm of the thalamus. Kybernetik 15, 27–37.
Lisman, J., Idiart, M., 1995. Storage of 7+/-2 short term memories in oscillatory subcycles. Science 267, 1512–1515.
Marreiros, A.C., Daunizeau, J., Kiebel, S.J., Friston, K.J., 2008. Population dynamics: variance and the sigmoid activation function. Neuroimage 42 (1), 147–157.
Mattia, M., Del Giudice, P., 2004. Finite-size dynamics of inhibitory and excitatory interacting spiking neurons. Phys. Rev., E Stat. Nonlinear Soft Matter Phys. 70, 052903.
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.
Moran, R.J., Kiebel, S.J., Stephan, K.E., Rombach, N., O'Connor, W.T., Murphy, K.J., Reilly, R.B., Friston, K.J., 2008. Bayesian estimation of synaptic physiology from the spectral responses of neural-masses. NeuroImage 42 (1), 272–284 Aug 1.
Mormann, F., Fell, J., Axmacher, N., Weber, B., Lehnertz, K., Elger, C., Fernandez, G., 2005. Phase/amplitude reset and theta-gamma interaction in the human medial temporal lobe during a continuous word recognition memory task. Hippocampus 15, 890–900.
Morris, C., Lecar, H., 1981. Voltage oscillations in the barnacle giant muscle fiber. Biophys. J. 35, 193–213.
Nykamp, D.Q., Tranchina, D., 2000. A population density approach that facilitates large-scale modeling of neural networks: analysis and an application to orientation tuning. J. Comput. Neurosci. 8, 19–50.
Omurtag, A., Knight, B.W., Sirovich, L., 2000. On the simulation of large populations of neurons. J. Comput. Neurosci. 8, 51–63.
Palva, J.M., Palva, S., Kaila, K., 2005. Phase synchrony among neuronal oscillations in the human cortex. J. Neurosci. 25, 3962–3972.
Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004. Comparing dynamic causal models. NeuroImage 22, 1157–1172.
Penny, W.D., Duzel, E., Miller, K.J., Ojemann, J.G., 2008. Testing for nested oscillation. J. Neurosci. Methods 174 (1), 50–61.
Risken, H., 1996. The Fokker–Planck equation. Springer, Berlin.
Robinson, P.A., Rennie, C.J., Rowe, D.L., O'Connor, S.C., Wright, J.J., Gordon, E., Whitehouse, R.W., 2003. Neurophysical modeling of brain dynamics. Neuropsychopharmacology 28 (Suppl. 1), S74–S79.
Rodriguez, R., Tuckwell, H.C., 1996. Statistical properties of stochastic nonlinear dynamical models of single spiking neurons and neural networks. Phys. Rev. E 54, 5585.
Rodriguez, R., Tuckwell, H.C., 1998. Noisy spiking neurons and networks: useful approximations for firing probabilities and global behavior. Biosystems 48 (1–3), 187–194.
Rodriguez, R., Tuckwell, H.C., 2000. A dynamical system for the approximate moments of nonlinear stochastic models of spiking neurons and networks. Math. Comput. Model. 31, 175–180.
Sirovich, L., 2003. Dynamics of neuronal populations: eigenfunction theory; some solvable cases. Network 14, 249–272.
Steriade, M., 2006. Grouping of brain rhythms in corticothalamic systems. Neuroscience 137, 1087–1106.
Tass, P.A., 2003. Stochastic phase resetting of stimulus-locked responses of two coupled oscillators: transient response clustering, synchronization, and desynchronization. Chaos 13, 364–376.
Tass, P., Rosenblum, M.G., Weule, J., Kurths, J., Pikovsky, A.S., Volkmann, J., Schnitzler, A., Freund, H.-J., 1998. Detection of n:m phase locking from noisy data: application to magnetoencephalography. Phys. Rev. Lett. 81, 3291–3294.
Tuckwell, H.C., Rodriguez, R., 1998. Analytical and simulation results for Stochastic Fitzhugh-Nagumo neuron and neural networks. J. Comput. Neurosci. 5, 91–113.
Vanhatalo, S., Palva, J.M., Holmes, M.D., Miller, J.W., Voipio, J., Kaila, K., 2004. Infraslow oscillations modulate excitability and interictal epileptic activity in the human cortex during sleep. Proc. Natl. Acad. Sci. U. S. A. 101, 5053–5057.
Valdes, P.A., Jimenez, J.C., Riera, J., Biscay, R., Ozaki, T., 1999. Nonlinear EEG analysis based on a neural-mass model. Biol. Cybern. 81, 415–424.
Varela, F., Lachaux, J.P., Rodriguez, E., Martinerie, J., 2001. The brainweb: phase synchronization and largescale integration. Nat. Rev., Neurosci. 2, 229–239.
Von Stein, A., Chiang, C., König, P., 2000. Top-down processing mediated by interareal synchronization. Proc. Natl. Acad. Sci. U. S. A. 97, 14748–14753.
Wendling, F., Bellanger, J.J., Bartolomei, F., Chauvel, P., 2000. Relevance of nonlinear lumped-parameter models in the analysis of depth- EEG epileptic signals. Biol. Cybern. 83, 367–378.
Wright, J.J., Rennie, C.J., Lees, G.J., Robinson, P.A., Bourke, P.D., Chapman, C.L., Gordon, E., Rowe, D.L., 2003. Simulated electrocortical activity at microscopic, mesoscopic, and global scales. Neuropsychopharmacology 28 (Suppl. 1), S80–S93.
Zetterberg, L.H., Kristiansson, L., Mossberg, K., 1978. Performance of a model for a local neuron population. Biol. Cybern. 31, 15–26.