Published in final edited form as:

J Stat Mech. ; 2013: P03003-. doi:10.1088/1742-5468/2013/03/P03003.

Beyond mean field theory: statistical field theory for neural networks

Michael A Buice¹ and Carson C Chow²

Michael A Buice: mabuice@mail.clm.utexas.edu; Carson C Chow: carsonc@niddk.nih.gov

¹Center for Learning and Memory, University of Texas at Austin, Austin, TX, USA

²Laboratory of Biological Modeling, NIDDK, NIH, Bethesda, MD, USA

Abstract

Mean field theories have been a stalwart for studying the dynamics of networks of coupled neurons. They are convenient because they are relatively simple and possible to analyze. However, classical mean field theory neglects the effects of fluctuations and correlations due to single neuron effects. Here, we consider various possible approaches for going beyond mean field theory and incorporating correlation effects. Statistical field theory methods, in particular the Doi–Peliti–Janssen formalism, are particularly useful in this regard.

Keywords

finite-size scaling; dynamics (theory); neuronal networks (theory); Boltzmann equation

1. Introduction

The brain is a complex system with a very large number of neurons that are individually complex and heterogeneous. Presently, we can either study neuronal dynamics with biophysical detail for a modest number of neurons or consider population effects in the mean field limit, where the contributions of individual neurons are averaged out. Mean field population rate or activity equations have been a standard tool of computational and theoretical neuroscience since the work of Wilson and Cowan, Cohen and Grossberg, Amari and Hopfield [1]–[5]. These models are phenomenological theories that obey certain properties in the vein of the Landau–Ginzburg Hamiltonian and are not formally derived from a microscopic theory.

Models of this type have been used to investigate pattern formation, visual hallucinations, content addressable memory and many other phenomena [6]–[9]. These equations describe the evolution of a neural activity variable often ascribed to the firing rate or synaptic drive of a population of interacting neurons [10, 11]. These equations are thought to represent the neural dynamics averaged over time or population of a more complicated underlying microscopic process. They capture the dynamics of the mean firing rate or activity that is independent of the influence of correlations, which in some cases may alter the dynamics considerably. As an example, the effects of synchrony, which have been proposed to be important for neural processing [12, 13], are not included.

What would be useful and desirable is to calculate the discrete neuron effects such as synchronous correlations or finite-size fluctuations associated with the mean rates. This would be the next level of complication beyond mean field theory. Better yet would be a set of self-consistent effective equations that carry the population activity dynamics as well as correlation information due to discrete neuronal fluctuations. Such a system could, for example, be used to analyze the dynamic effects of correlation-based learning algorithms such as spike-timing-dependent plasticity, which cannot be analyzed in mean field theory.

There are two possible approaches towards this end—top down or bottom up. In the top down approach, one could posit a microscopic theory consistent with mean field theory and then derive the correlation equations for that theory. In the bottom up approach, one could start with a biophysically motivated spiking neural network and then derive an effective theory from that by marginalizing or tracing over microscopic degrees of freedom. Both have advantages and disadvantages. The top down approach is more straightforward but not guaranteed to be unique. The bottom up approach is explicit but much more difficult to calculate.

2. Mean field rate equations

A paradigmatic neural mean field theory is the Wilson–Cowan equation:

$$\dot{a}_i(t) = -\alpha a_i(t) + f \left(\sum_j w_{ij} a_j + I_i(t) \right) \quad (1)$$

where $a_i(t)$ is the time-dependent activity of the network at location i , f is a gain function, w_{ij} is a coupling kernel between neurons at different locations and $I_i(t)$ is an external input. The classic derivation of the Wilson–Cowan equation is phenomenological in the Landau–Ginzburg sense. The interpretation of the activity field is not specified. It is usually thought of as a mean neuronal firing rate, but it is not clear whether it is the rate averaged over a pool of neurons or a time average. It is also not clear what sets the time constant α^{-1} . It may be due to some microscopic parameter, such as the membrane time constant, or some collective property. The same issue holds for the gain function. The Wilson–Cowan equation is often coarse grained onto a continuous configuration variable x

$$\partial_t a(x, t) = -\alpha a(x, t) + f \left(\int_{\Omega} d\vec{y} w(x, y) a(y, t) + I(x, t) \right). \quad (2)$$

The heuristic derivation of these equations overlooks the underlying dynamics and the ensuing statistics. In order to address effects in these systems which are ‘beyond’ mean field theory, one is thus obligated to specify the underlying dynamics, or in some way constrain the space of possible extensions (or rather definitions) of the model. In the following two sections, we describe two approaches. In the first, we specify various possible underlying stochastic models whose mean field descriptions match a form of the Wilson–Cowan equations. This approach has the advantage that the rate equations arise ‘naturally’ and that the beyond mean field effects are straightforwardly computed. In section 4, we will explore approaches which build statistical descriptions of networks from detailed underlying models

in which the statistical behavior arises through network heterogeneity rather than intrinsic stochasticity.

3. Stochastic microscopic dynamics

In this section we construct stochastic models underlying the rate equations. To begin, let us consider what might be the ‘simplest’ model which obeys the given rate equation. Following the heuristic description given in Wilson–Cowan [1, 2], we interpret the activity $a_i(t)$ as the probability density of a neuron being ‘active’ at time t . This requires that we consider neurons to have an ‘active’ state and a ‘quiescent’ state. For each cortical region i , the state of the system is thus described by the number of neurons $n_i(t)$ which are in the ‘active’ state at time t . We consider a model appropriate for a low firing rate regime and so we formally regard the total number of neurons in each region to be infinite. The stochastic dynamics must be defined in terms of transitions between these states $n_i(t)$. Using the Wilson–Cowan rate equation, we suppose that there are transitions $n_i \rightarrow n_i - 1$ (an ‘active’ neuron ‘decays’, i.e. becomes inactive) at rate α and $n_i(t) \rightarrow n_i(t) + 1$ at rate $F_i[n(\vec{t})]$ (corresponding to activation). Here F_i is dependent upon the number of neurons active in the network. This is depicted in figure 1.

Naturally, this system is described by a probability distribution $P(n, \vec{t})$ which evolves according to the transition probabilities given above. This gives the following master equation, considered in [14] and [15]

$$\frac{dP(\vec{n}, t)}{dt} = \sum_i [\alpha(n_i + 1)P(\vec{n}_{i+}, t) - \alpha n_i P(\vec{n}, t) + F_i(\vec{n}_{i-})P(\vec{n}_{i-}, t) - F_i(\vec{n})P(\vec{n}, t)] \quad (3)$$

where $P(n, \vec{t})$ is the probability of the network having the activity configuration described by n at time t . The local neural activity at location i is characterized by the number of active neurons n_i . Each neuron is active with a characteristic time scale of $1/\alpha$. Configurations $\vec{n}_{i+}$ and $\vec{n}_{i-}$ denote the configuration $\vec{n}$ where the i th component is $n_i \pm 1$, respectively. The rate at which a neuron at location i becomes active is given by the firing rate or gain function $F_i(\vec{n})$, which is an implicit function of the coupling weight function w_{ij} and external inputs I_i . We specifically use a different symbol for the gain function $F_i(\vec{n})$, which appears in (3) because it may not be the same function as the mean field gain $f(I_i(t) + \sum_j w_{ij} n_j)$, which appears in (1).

The most direct and obvious method of analyzing the master equation (3) is to consider the moment hierarchy, which is derived by taking expectation values of the master equation with powers of $n_i(t)$. Naturally, if F_i is a nonlinear function of $n(\vec{t})$, then this moment hierarchy will not consist of a finite closed set of equations. A natural assumption at this point is to consider that the higher moments (or more precisely the cumulants) are suppressed relative to the lower moments, in particular that the mean is dominant. Using this approximation, the lowest level of the hierarchy is given by

$$\partial_t a_i(t) = -\alpha a_i(t) + F_i(\vec{a}) + \text{h.o.t.} \quad (4)$$

where $\langle n_i(t) \rangle = a_i(t)$ and the higher order terms (h.o.t.) are assumed small. This is a small noise approximation which assumes that the spike activity is largely deterministic, with only minimal fluctuations around the mean. Neglecting the higher order terms produces *mean field theory*. Bressloff [16] demonstrates an analysis of such an expansion in which the system size is used as a truncation parameter. In that case, the master equation would be reinterpreted for a population with a finite-size N , with appropriate boundary conditions, and the system self-averages over a large population, which suppresses the population level fluctuations, making mean field theory valid. In the system size expansion, the second moment of the fluctuations scale as the inverse system size and vanish in the infinite system size mean field limit. The mean field gain function is the same as that in the master equation.

However, the heuristic argument of the Wilson–Cowan analysis is that the activity represents a probability of firing. We may expect solutions of the master equation in that case to be approximations to a multi-component Poisson distribution. The Poisson distribution has the property that the mean is the same as the variance, so that the system size expansion, where noise is suppressed in the mean field limit, is not appropriate. We can instead derive our moment hierarchy for *normal ordered moments* or *factorial moments*. For a single variable, the m th factorial moment is given by $\langle n!/(n-m)! \rangle$ instead of $\langle n^m \rangle$. In a factorial moment the Poisson terms have been ‘subtracted’ from each moment (the factorial moments of a Poisson distribution other than the mean are all zero). The master equation gain function F_i is an explicit function of $n_i(t)$. The expectation value can thus be expressed as an expansion in terms of the moments of $n_i(t)$. However, in order to develop a self-consistent factorial moment hierarchy, we need to reorganize the expansion of the gain function in terms of factorial moments, which produces a new gain function f . The lowest order equation of the hierarchy in terms of factorial moments is given by

$$\partial_t a_i(t) = -\alpha a_i(t) + f_i(\vec{a}) + \text{h.o.t.} \quad (5)$$

The approximation which ignores the higher order terms constitutes a solution of the master equation (3) which is approximately a Poisson distribution with mean given by $a_i(t)$ determined by equation (5). Because of this, we claim that the master equation (3) can be viewed as a ‘minimal’ extension of the Wilson–Cowan equations. As pointed out in Bressloff [16], the choice of master equation leading to a particular Wilson–Cowan equation is not unique, so one cannot claim via this construction that the biophysical information leading to mean field theory is sufficient to determine the correlation structure. The mean of a distribution is not sufficient to determine the entire distribution. The fluctuation terms derived from the hierarchy are thus natural candidates for ‘extending’ mean field theory. We can compute these from the moment hierarchy, in a similar manner to those derived in the system size expansion, which was done explicitly in [15]. Note that the equations at second order do not agree. The small noise expansion is different from the near Poisson expansion.

Alternatively, we can derive these moment equations using a formalism developed in the context of reaction–diffusion systems in physics [17]–[20]. Although somewhat abstract and

technical, this method expresses the solutions to the master equation in terms of variables $\phi_i(t)$, $\tilde{\phi}_i(t)$ which naturally represent Poisson solutions. When correlations between the variables $\phi_i(t)$ for different i and t are zero, the solution to the master equation is given by a Poisson distribution with mean $\langle \psi_i(t) \rangle$. The higher moments describe deviations from Poisson behavior. This has the advantage of making the near Poisson approximation more 'obvious' as a result of the nature of the counting process underlying the master equation and providing some quantifiable description of the exact nature of the approximation made in truncating the hierarchy. In addition, the formalism has other advantages for computing beyond mean field effects and for relating the master equation to other statistical models which may underly the Wilson–Cowan equations, as we will see. This alternative approach [14, 15] considers the moment generating functional for the master equation, $Z[J_i(t), \tilde{J}_i(t)]$. Derivatives of the generating functional evaluated at $J_i(t) = \tilde{J}_i(t) = 0$ provide the moments of the solution of the master equation, including unequal time correlation functions (which would not be captured by a standard generating function of the distribution at a single time point). The generating functional is given by the form

$$Z[J_i(t), \tilde{J}_i(t)] = \int \mathcal{D}\tilde{\varphi} \mathcal{D}\varphi e^{-S[\tilde{\varphi}, \varphi]} e^{\sum_i \int (\tilde{\varphi}_i(t) J_i(t) + \varphi_i \tilde{J}_i(t)) dt} \quad (6)$$

where

$$S[\tilde{\varphi}, \varphi] = \frac{1}{h} \sum_i \int dt \left[\tilde{\varphi}_i \frac{\partial}{\partial t} \varphi_i + \alpha \tilde{\varphi}_i \varphi_i - \tilde{\varphi}_i f \left(\sum_j w_{ij} (\tilde{\varphi}_j \varphi_j + \varphi_j) \right) \right] - \sum_i W[\tilde{\varphi}_i(0)] \quad (7)$$

is the *action*, a quantity equivalent to the negative log likelihood for the complex variables ϕ_i and $\tilde{\phi}_i$, which are related to the spike counts through $n_i = \phi_i \tilde{\phi}_i + \phi_i$, h is a scaling factor, and $W[\tilde{\varphi}_i(0)]$ describes the initial distribution. For example, if the initial state is described by Poisson statistics, $W[\tilde{\varphi}_i(0)] = a_i^0 \tilde{\varphi}_i(0)$, where a_i^0 is the mean of the Poisson distribution at i . The activity variable in (1) is given by the expectation value $\langle \phi_i \rangle$ and the factorial variance (i.e. deviation of the variance from the Poisson value) is given by $\langle \phi_i \phi_j \rangle$. The higher order moments of ϕ_i are factorial moments of the spike count n_i so that they are identically zero for Poisson firing neurons. A functional integral is an integral over the domain of functions, in this case the complex functions $\phi_i(t)$ and $\tilde{\phi}_i(t)$. The integrations run over the real axis for $\phi_i(t)$ and the imaginary axis for $\tilde{\phi}_i(t)$, for all t . It is important to note that the action (7) is completely equivalent to the master equation (3). It does not represent any form of approximation, it is simply another form of expressing the master equation. The parameter h allows us to organize the expansion in fluctuations, which provides a natural way of truncating the moment hierarchy. This expansion produces the factorial moment expansion without any extra manipulation; it is a by-product of the expression of the master equation as a counting process and the Doi–Peliti path integral formalism.

The derivation of the action follows from recasting the master equation in terms of creation and annihilation operators and then transforming to a coherent states representation [14, 15]. Heuristically, one can consider the master equation to be equivalent to an effective Langevin

equation for the spike counts with a fictitious stochastic forcing function that is consistent with the master equation. The probability density function in the generating functional (6) can then be written formally in terms of the path integral

$$P(\vec{n}, t) \propto \sum_i \int \mathcal{D}\xi_i(t) \delta[\dot{n}_i + \alpha n_i - F(\vec{n}) + \xi_i(t)] P[\xi_i(t)] \quad (8)$$

where $\delta[\cdot]$ is a Dirac delta functional enforcing the counts to obey the effective Langevin equation. $P[\xi_i(t)]$ is the probability density functional of the fictitious noise source $\xi_i(t)$ that obeys

$$\sum_j \int \mathcal{D}\xi_j(t) e^{\sum_i \int \tilde{n}_i(t) \xi_i(t) dt} P[\xi_j(t)] = \exp[(e^{\tilde{n}} - 1 - \tilde{n})F + \alpha n(e^{\tilde{n}} - 1 + \tilde{n})] \quad (9)$$

which is the generating functional for Poisson activation with rate F and decay with rate α ; as required. The transformation $\phi_i = n_i \exp(-\tilde{n}_i)$, $\tilde{\phi}_i = \exp(\tilde{n}_i) - 1$, effectively ‘linearizes’ the Poisson noise, resulting in the action (7). Restoring normal ordering accounts for the transition from F to f . The moments of ϕ_i are factorial moments of the counts.

More formally, the activity a_i , defined as the expectation value of ϕ_i , is obtained by taking derivatives of the generating functional $Z[J, \tilde{J}]$, with respect to J_i . An activity equation can be found by Legendre transforming $\ln Z[J, \tilde{J}]$ to an ‘effective action’ $\Gamma[a(x, t)]$. The extremum of the effective action gives the equation of motion for $a_i(t)$, which is the Wilson–Cowan equation. To derive coupled equations for the activity $a_i(t)$ and the correlations $C_{ij}(t) = \langle \phi_i \phi_j \rangle$, we augment the generating functional by adding a term proportional to $K_{ij} \tilde{\phi}_i \phi_j + \dots$ and then Legendre transforming in all the variables J and K to obtain the generalized effective action from which equations of motion for these composite operators can be calculated [21]. This can be continued for an arbitrary number of moments.

The resulting generalized equations depend on the interpretation of the expansion parameter h . If h is taken to be unity, as was chosen in [15], then the corresponding expansion parameter is the inverse distance from criticality. To tree level, the result is that the equation for $a(x, t)$ is [15]:

$$\partial_t a_i(t) = -\alpha a_i(t) + f \left(\sum_j w_{ij} a_j(t) + I_i(t) \right) + \frac{1}{2} f'' \left(\sum_j w_{ij} a_j(t) + I_i(t) \right) \sum_{jk} w_{ij} w_{ik} C_{jk}(t) \quad (10)$$

$$\frac{d}{dt} C_{ij}(t) = -2\alpha C_{ij}(t) + f'(s_i) \sum_k w_{ik} C_{kj}(t) + f'(s_j) \sum_k w_{jk} C_{ki}(t) + f'(s_i) w_{ij} a_j(t) + f'(s_j) w_{ji} a_i(t). \quad (11)$$

The loop expansion of the effective action will produce arbitrary corrections due to the higher order effects described in the action. By contrast, the system size expansion from Bressloff [16] has the same form, but the source term for the C_{ij} equation is replaced by

$h[aa_i + f(s_i)]\delta_{ij}$ and $1/N$ serves as the value of h . A major advantage of the path integral formalism over a direct analysis of the master equation or moment hierarchy is that diagrammatic methods of field theory developed over the past half century are straightforwardly applied for all types of perturbative calculations. Compare the arduous and delicate calculations required to derive the generalized activity equations (10) and (11) directly from the master equation with the path integral approach [15]. The system size expansion of [16] could also be computed using similar machinery by employing a loop expansion in the inverse system size. Semi-classical WKB approximations of large deviation theory can also be directly implemented in terms of path integrals.

One of the primary uses of the path integral formalism in physics is the analysis of scaling properties of systems. For the present model it is shown in Buice and Cowan [14] that, assuming that the firing rate function f is a ‘saturating’ function, the system described by the master equation (3) displays the critical behavior of a model called *directed percolation*, which describes activity of phase transitions of a scalar order parameter away from a non-fluctuating absorbing state.

Another advantage of the path integral formalism is the ease which it affords us in connecting to different models. The Markov system described by (3) is a stochastic process where the probability of the spike count at a given location is increased by counts in other locations and decreases at a fixed rate. This may be construed to be a contrived abstract process. Let us introduce a class of models, first introduced by Cowan [22], in which neurons are considered to be in one of q fixed states. A reasonable set of states to consider is the triplet of *active*, *quiescent*, and *refractory*. Figure 2 depicts a pair of such models with transitions between the states, where the neurons can be in one of either two or three states. The state of the network is then given by a vector v whose i th component gives the state of neuron i . The probability distribution is gives the probability for the network to be in a collective state v . Readers familiar with the Forest Fire model [23] will note the similarity with the three-state Cowan models.

The master equations for these models are more simply expressed in the algebraic representations necessary for formulating the generating functional. The algebraic construction as well as the derivation of the actions appears in [24]. Applying the Doi–Peliti formalism to the algebraic representation gives the following action for the two-state model:

$$\begin{aligned} \mathcal{S}(\tilde{\psi}, \psi, \tilde{\phi}, \phi) = & \sum_i \int_{-\infty}^{\infty} dt \left[\tilde{\psi}_i(t) \partial_t \psi_i(t) + \tilde{\phi}_i(t) \partial_t \phi_i(t) + \alpha \left(|\psi_i(t)|^2 - \tilde{\phi}_i(t) \psi_i(t) \right) \right. \\ & + \left. \left(|\phi_i(t)|^2 - \tilde{\psi}_i(t) \phi_i(t) \right) f \left(\sum_j w_{ij} \{ |\psi_j(t)|^2 + \psi_j(t) \} \right) \right] \\ & + \sum_i \left(p_i \tilde{\psi}_i(0) + q_i \tilde{\phi}_i(0) \right) \end{aligned} \quad (12)$$

where the variables $(\tilde{\psi}_i, \psi_i)$ and $(\tilde{\phi}_i, \phi_i)$ correspond to the active and quiescent states, respectively, and p_i, q_i represent the initial probabilities of neuron i being in the active or quiescent state. The expectation values $\langle \psi_i(t) \rangle, \langle \phi_i(t) \rangle$ give the *probabilities* of neuron i being active or quiescent. Necessarily $\langle \psi_i(t) \rangle + \langle \phi_i(t) \rangle = p_i + q_i = 1$. The higher moments give *joint* probabilities. The three-state action is similar, with the inclusion of variables $\pi_i(t)$,

$\pi_i(t)$ giving the probability of neuron i being refractory. Despite the differences in the models and state space descriptions, note the similarities between (7) and (12). We can apply the same techniques in order to derive mean field equations and corrections. The primary difference in this case is the natural saturation effect which arises as a result of the quiescent state. Defining $p_i(t) = \langle \psi_i(t) \rangle$, $q_i(t) = \langle \varphi_i(t) \rangle$, the mean field equation for the two-state model is:

$$\dot{p}_i(t) + \alpha p_i(t) = (1 - p_i(t)) f \left(\sum_j w_{ij} p_j(t) \right) \quad (13)$$

where we have used $p_i(t) + q_i(t) = 1$. Note that when the probability of neuron i being in the active state reaches 1, there is no further input from the firing rate function f .

An important aspect of the Cowan models is that the transition rates are dependent at most upon the number of active neurons at time t . This means that the corresponding master equation or, similarly, the action, for those variables is exactly solvable. Indeed, we can see this in the form of the actions. The non-active state variables are bilinear in tilde and non-tilde variables, which means we can formally solve for the functional integral in terms of the active state variables, leaving an action defined over only the remaining active state. This is given by

$$\mathcal{S}(\tilde{\psi}, \psi) = \sum_i \int_{-\infty}^{\infty} dt \left[\tilde{\psi}_i(t) \partial_t \psi_i(t) + \alpha |\psi_i(t)|^2 - \tilde{\psi}_i(t) \Phi_i(t) f \left(\sum_j w_{ij} \{ |\psi_j(t)|^2 + \psi_j(t) \} \right) \right] + \sum_i p_i \tilde{\psi}_i(0) \quad (14)$$

where $\Phi_i(t)$ is the ‘classical’ solution for the quiescent state operator $\varphi_i(t)$, which will in general be a function of the fields $\tilde{\psi}_i(t)$, $\psi_i(t)$. This action differs from (7) only in the saturating term. If we assume that the probability of being in the active state is small, i.e. we assume that the network is in a low firing rate regime, then $\Phi_i(t) \approx 1$ (which also neglects the non-Markovian properties introduced by the marginalization). Thus we see that the action for low firing rates is given by the action for the effective spike model. As a consequence, the two models will share common critical properties, in particular both will be members of the directed percolation universality class. This is verified for the Forest Fire model in [25]. We see that the action (7) is quite general and can describe a variety of stochastic dynamics.

4. Deterministic microscopic dynamics

The stochastic approach begins with plausible probabilistic microscopic dynamics, or postulates one that is consistent with the Wilson–Cowan equation. As we showed in section 3, the resultant action (7) from a simple probabilistic model is quite general. The fact that a class of more general models (e.g. the 2- and 3-state Cowan models) reduces to an equivalent model gives us some hope that a similar reduction may hold for more ‘realistic’ neural models, say those described by a set of Hodgkin–Huxley neurons. These neuron models generally consist of a system of deterministic ordinary differential equations. (Our formalism can also be generalized to include stochastic forcing.) The orbit of a neuron is

completely specified by the initial conditions and parameter values. However, uncertainty is introduced if we consider the family of orbits over a heterogeneous set of initial conditions or parameters. The heterogeneity imposes a distribution over the space of states and gives rise to dynamics which appear ‘noisy’. A simple example of a model exhibiting this kind of behavior is an ideal gas, which obeys completely deterministic dynamics, but which shows apparently stochastic behavior in its constituent particles. As in the case of a gas, kinetic theory may be applied to analyze the dynamics of neural systems. We demonstrate here how a kinetic theory approach informs us about effects beyond mean field theory. It should be noted that specifying initial or parametric noise is not always sufficient to insure that a system will exhibit dynamical noise in the stationary limit (or even that it has a stationary limit). Our method, however, will still admit the analysis of such systems. It should also be noted that whereas in section 3 we were concerned with the analysis of a system given its underlying statistics, in the present section we add the additional question of how such statistics may arise, or indeed whether they remain (e.g. noise in the initial conditions is not guaranteed to remain in the stationary limit. The models we analyze are all strictly deterministic, with disorder arising only in the distribution of initial conditions and network parameters.

Consider a network of interacting neurons with the general form

$$\frac{dV_i}{dt} = I_i(\vec{z}_i, V_i, t) + \sum_j g_{ij}[V_i, V_j, t] \quad (15)$$

where V_i is the membrane potential of neuron i , $\vec{z}_i$ are a set of neuron-specific time-dependent internal variables, and g_{ij} is a synaptic input from neuron j to neuron i . Generally, g_{ij} is a time-dependent function that is triggered when the membrane potential crosses a threshold and decays with a specific time course depending on the synaptic and neurotransmitter type.

In many cases, the dynamics of a single neuron can be projected onto a phase variable [26, 27], resulting in a set of N neuron equations obeying

$$\dot{\theta}_i = F_i(\vec{\theta}, t) \quad (16)$$

where θ_i is the phase of neuron i , $\vec{\theta} = \{\theta_1, \theta_2, \dots, \theta_N\}$, and $F_i(\vec{\theta}, t)$ is the phase and time-dependent coupling for neuron i , which depends explicitly on the other neuron phases and implicitly on some set of fixed parameters. Note, that the neurons need not be weakly coupled for a phase formulation.

Equation (16) describes a set of deterministic differential equations. The apparent stochastic behavior appears upon examining the population level behavior, or upon comparing the behavior of a single individual element to the larger population. In order to characterize the population level behavior, we define the population density

$$\eta(\theta, \gamma, t) = \frac{1}{N} \sum_{i=1}^N \delta(\theta - \theta_i(t)) \delta(\omega - \omega_i) \quad (17)$$

where $\delta(\cdot)$ is the point mass or Dirac delta functional, $\theta_i(t)$ are the solutions to system (16), and ω_i is a parameter (or set of parameters) associated with oscillator i . Oscillators are characterized by both their phase and parameter values. $\eta(\theta, \omega, t)$ is essentially the count (more precisely, the fraction) of the number of neurons with phase θ and parameters ω at time t .

The population density obeys the conservation equation

$$\frac{\partial}{\partial t} \eta + \frac{\partial}{\partial \theta} F \eta = 0 \quad (18)$$

which is known as the Klimontovich equation in kinetic theory and is valid only in the weak or distributional sense since η is not differentiable. The Klimontovich equation fully specifies the time-dependent evolution of the system and is entirely equivalent to the original differential equation. However, it is difficult to work with due to the weak nature of the solution. The reason for its introduction is that we are interested in the statistics of the population. Each realization of the network over the distributions of the initial conditions and the network and neuron parameters generates a realization of η . The population statistics are naturally described by the distribution of η induced by the distributions of initial conditions and network parameters. We can characterize this distribution by taking expectation values of the continuity equation and constructing a moment hierarchy, in the same way we constructed a hierarchy for the master equation (3).

Averaging over an ensemble of similarly prepared systems with different initial conditions and parameters and denoting averages over this measure by $\langle \cdot \rangle$, the average of (18) yields the equation

$$\frac{\partial}{\partial t} \langle \eta \rangle + \frac{\partial}{\partial \theta} \langle F \eta \rangle = 0. \quad (19)$$

Since F is an implicit function of η , the average of $F \eta$ will depend on higher order moments. The moments in this hierarchy are $\langle \rho_n \rangle \equiv \langle \eta^n \rangle$ and are called n -neuron distribution functions. These represent the joint probability distribution for finding a set of neurons with the values of θ, ω at time t , where the probability necessarily refers to the incidence in the ensemble. Each of the moments will in turn depend on even higher order moments, leading to a moment hierarchy. This hierarchy is similar, but not identical, to the moment hierarchy in section 3. In that case, the moments in the hierarchy were for the activity. However, in this case, the moments are in terms of the population density. Needless to say, even truncations of the population density moment hierarchy are extremely difficult to solve.

As before, mean field theory is given by neglecting all correlations and higher order cumulants, leading to an approximate equation for the 1-neuron distribution function. This is equivalent to setting $\langle F \rangle = \langle F \rangle$, which gives the self-consistent mean field system

$$\frac{\partial}{\partial t} \rho = - \frac{\partial}{\partial \theta} F(\theta, \omega) \rho$$

where $\rho_1(\theta, \omega, t) = \langle \eta(\theta, \omega, t) \rangle \equiv \rho$. If we specify the phase $\theta = \pi$ as ‘firing’, then the number of firing neurons is characterized by $\rho(\pi, \omega, t)$. The flux at π , $\partial \rho(\theta, \omega, t) / \partial \theta|_{\theta=\pi}$, is the firing rate. Higher order density moments ρ_n are likewise defined. We will show how all moments can be computed using a path integral approach.

To take a concrete example, consider the Kuramoto model

$$\dot{\theta}_i = \omega_i + \frac{K}{N} \sum_j f(\theta_j - \theta_i). \quad (20)$$

This is the weak coupling limit of a network of sparsely coupled integrate-and-fire neurons [28]. The Klimontovich equation is given by

$$\mathcal{C}(\eta) \equiv \frac{\partial \eta}{\partial t} + \omega \frac{\partial \eta}{\partial \theta} + K \frac{\partial}{\partial \theta} \int_{-\infty}^{\infty} \int_0^{2\pi} f(\theta' - \theta) \eta(\theta', \omega', t) \eta(\theta, \omega, t) d\theta' d\omega' = 0. \quad (21)$$

Hildebrand *et al* [29] explicitly constructed and truncated the moment hierarchy from the Klimontovich equation. Buice and Chow then constructed an equivalent Doi–Peliti field theory [17]–[20], [30]. Since the model is deterministic, the deterministic time evolution of $\eta(\theta, \omega, t)$ serves to map the initial distribution forward in time. The functional probability measure $\mathcal{P}[\eta(\theta, \omega, t)]$ for the density $\eta(\theta, \omega, t)$ is a delta functional that enforces the deterministic evolution from equation (21) along with an expectation taken over the distribution $\mathcal{P}_0[\eta_0]$ of the initial configuration $\eta_0(\theta, \omega) = \eta(\theta, \omega, t_0)$. Hence we arrive at the following path integral

$$\mathcal{P}[\eta(\theta, \omega, t)] = \int \mathcal{D}\tilde{\eta} \mathcal{D}\eta_0 \mathcal{P}_0[\eta_0] \exp(-N \int d\theta d\omega dt \tilde{\eta} [\mathcal{C}(\eta) - \delta(t - t_0) \eta_0(\theta, \omega)]) \quad (22)$$

$$\equiv \int \mathcal{D}\tilde{\eta} e^{-NS[\tilde{\eta}, \eta]} \quad (23)$$

where $\tilde{\eta}(\theta, \omega, t)$ is usually called the ‘response field’ after Martin–Siggia–Rose [31] and the integral is taken along the imaginary axis. We emphasize that no external dynamical noise is added to our system. Any statistical uncertainty completely derives from the distribution of the initial state and the disorder in the network and neural parameters. For simplicity, we assume that the initial phases and driving frequencies for each of the N oscillators are independent and obey the distribution $\rho_0(\theta, \omega)$, which leads to

$$S[\eta, \tilde{\eta}] = -N \int d\theta d\omega dt \tilde{\eta} \mathcal{C}(\eta) + N \ln \left[1 + \int d\theta d\omega \left[e^{\tilde{\eta}(\theta, \omega, t_0)} - 1 \right] \rho_0(\theta, \omega) \right]. \quad (24)$$

We see that the fluctuations (i.e. terms nonlinear in $\tilde{\eta}$) appear only in the initial condition of (24), which is to be expected since the Kuramoto system is deterministic. In this form the continuity equation (21) appears as a Langevin equation sourced by the noise from the initial state. The ‘noise’ is entirely determined by ‘boundary’ terms.

Although the stochasticity is inherited entirely from the initial conditions and parameters, it is not obvious how (or at this point, if) this results in dynamical ‘noise’. We can simplify the structure of the noise in the action (24) by performing a Doi–Peliti–Janssen transformation [20]:

$$\begin{aligned} \varphi(\theta, \omega, t) &= \eta \exp(-\tilde{\eta}) \\ \tilde{\varphi}(\theta, \omega, t) &= \exp(\tilde{\eta}) - 1. \end{aligned} \quad (25)$$

Under the transformation (25), the action (24) becomes

$$S[\varphi, \tilde{\varphi}] = \int d\omega d\theta dt \left[\tilde{\varphi} \left(\frac{\partial}{\partial t} + \omega \frac{\partial}{\partial \theta} \right) \varphi + K \int d\omega' d\theta' (\tilde{\varphi}' \tilde{\varphi} + \tilde{\varphi}) \frac{\partial}{\partial \theta} \left\{ f(\theta' - \theta) \varphi' \varphi \right\} \right] - \ln \left[1 + \int d\theta d\omega \tilde{\varphi}(\theta, \omega, t_0) \rho_0(\theta, \omega) \right]. \quad (26)$$

Correlation functions and higher moments can be computed for this action using a loop expansion around mean field theory, i.e. a smooth solution $\rho(\theta, \omega, t)$ of the continuity equation (21) with initial condition $\rho(\theta, \omega, t_0) = \rho_0(\theta, \omega)$. In general, a solution cannot be expressed in analytic closed form. However, explicit expressions can be computed for fluctuations around the stationary incoherent state $\rho(\theta, \omega, t) = \rho_0(\theta, \omega) = g(\omega)/2\pi$, where $g(\omega)$ is a fixed frequency distribution. The incoherent state is an exact solution of the continuity equation (21). If we transform the field variables via $\phi \rightarrow \phi + \rho$ in (26), the resulting shifted action with $\rho = g(\omega)/2\pi$ describes fluctuations about the true mean of the distribution. We can evaluate the moments of the probability distribution (23) with (26) using the loop expansion (i.e. method of steepest descents), in which $1/N$ appears as an expansion parameter [30, 32]. The vertex diagrams for (26) (shifted about the incoherent state) are given in figure 3.

An immediate result from the form of the action (and the loop expansion) is that in the infinite-size limit, all moments higher than the first are zero (provided the terms in the series are not at a singular point, i.e. the onset of synchrony). This is a consequence of the dependence of the action on N . Solutions to mean field theory (21) are given by the sum of those diagrams from the shifted action (26) which contribute to the mean at tree level (0 loops, or $O(1)$ in the $1/N$ expansion). If the initial conditions are smooth, then mean field theory is given by the relevant smooth solution of (21). In most of the previous work (e.g. [33, 34]), smooth solutions to (21) were taken as the starting point and hence automatically assumed mean field theory.

We have demonstrated how effective stochastic dynamics can be generated from uncertainty in the initial conditions and parameters. However, all systems need not obey this property.

For example, if a system had a strongly attracting fixed point, then all initial conditions would flow to the fixed point and the higher moments would all vanish in time. This situation is admissible even in the case of heterogeneity in the network parameters. Both of these extreme situations would be immediately detected by our formalism, e.g. had we considered the Kuramoto model in the synchronized regime with a distribution $g(\omega)$ with compact support. For a sufficiently large K , the entire ensemble distribution reduces to a delta functional on the ‘locked’ state.

One of Kuramoto’s goals in analyzing the mean field of his model was to develop an effective equation for the order parameter magnitude r , defined by

$$Z = \frac{1}{N} \sum_j e^{j\theta_j(t)} \equiv r(t) e^{i\psi(t)} \quad (27)$$

which is a measure of the number of synchronous oscillators in the network. In the mean field limit, there is a continuous phase transition from $r = 0$ to some $r > 0$ for some $K > K_c$. An actual finite-size simulation reveals, however, that r never quite reaches zero and fluctuates around some small value. Using our formalism, we computed these fluctuations in r , in terms of the 2-oscillator distribution function, as

$$\langle Z \bar{Z} \rangle = \langle r^2(t) \rangle = \int d\omega d\omega' d\theta d\theta' \langle \eta(\theta, \omega, t) \eta(\theta', \omega', t) \rangle e^{i(\theta - \theta')}. \quad (28)$$

The complete computation of the 2-oscillator distribution function can be found in [30], but the diagrams which contribute to it are shown in figure 4. The terms here are due only to the diagram in figure 4(a). This has a simple analytic form when $g(\omega)$ is chosen to be the Cauchy distribution ($g(\omega) = (1/\pi)(\gamma/(\gamma^2 + \omega^2))$).

$$\langle r^2(t) \rangle = \frac{1}{N} \frac{2\gamma}{2\gamma - K} - \frac{1}{N} \frac{K}{2\gamma - K} e^{-(2\gamma - K)t}. \quad (29)$$

At $K = 0$, this reduces to $\langle r^2 \rangle = 1/N$, which is the sampling error in the order parameter for a finite set of N neurons. The onset of synchrony is at $K = 2\gamma$. Note that (29) diverges as K nears this value. This is a standard symptom of these expansions, that each term in the loop expansion diverges near a bifurcation point.

Another outstanding problem left by Kuramoto was to determine the stability of the incoherent state, which was apparent from simulations but not from analysis. Strogatz and Mirollo [33] analyzed the mean field theory of the Kuramoto model to study this question and found that, in the absence of externally imposed noise, the mean field theory was actually marginally stable. A later paper [34] amended this result to describe how the stability of the order parameter Z was a manifestation of the dephasing effects in Landau damping, but this did not resolve the fundamental issue of the stability of mean field theory. To see this, consider a non-smooth perturbation in which, for example, a particular oscillator is given a particular spatial phase and driving frequency at some time t' . Mean field theory

predicts that the order parameter will *never* relax back to zero in response to this perturbation. In fact, as shown in [30], if the perturbed oscillator has driving frequency ω_0 and initial phase θ_0 , the perturbed order parameter evolves as

$$\delta\langle Z(t) \rangle = \frac{e^{i\theta_0}}{N} \frac{1}{\omega_0^2 + (\gamma - K/2)^2} \left[\left(\gamma \left(\gamma - \frac{K}{2} \right) + \omega_0^2 - \frac{K}{2} i \omega_0 \right) e^{-i\omega_0 t} - \left(-i\omega_0 + \gamma - \frac{K}{2} \right) \frac{K}{2} e^{-(\gamma - K/2)t} \right]. \quad (30)$$

Note that this never relaxes to zero and will always have a complex component in the direction of the perturbed oscillator $\theta_0 - \omega_0 t$. This marginal stability manifested in the Strogatz–Mirollo analysis as a continuous spectrum of the linear response obeying $s + in\omega_0 = 0$, where s represents the element of the spectrum. From the path integral one can compute corrections to the stability due to the finite-size effects, i.e. compute the corrections to the spectrum. This has the effect of altering the spectrum to be the zeros of

$$s + in(\omega + \delta\omega) + n^2 D = 0 \quad (31)$$

where

$$\delta\omega = -\frac{K^2}{2N} \frac{\omega}{(\gamma - K/2)^2 + \omega^2} \left[\frac{4\gamma - K}{2\gamma - K} \right] \quad (32)$$

$$D = \frac{K^2}{2N} \frac{\gamma}{(\gamma - K/2)^2 + \omega^2}. \quad (33)$$

Note that there is a shift in driving frequency and that it is $O(1/N)$ in magnitude. Note also that there appears an operator D which is proportional to n^2 . This is equivalent to an externally supplied diffusion, where the diffusion constant depends upon the network and oscillator parameters.

We have thus demonstrated how incomplete information about the state of a fully deterministic system leads to a natural probabilistic description. We now extend this to the apparent stochastic behavior of network elements, i.e. the individual neurons. As we discussed above, the distribution over initial conditions and network parameters imposes a distribution over realizations of the network, whose moments are given by $\langle \eta^n \rangle$. In a similar way, there is also an imposed distribution over the trajectories of single oscillators in the network, which will depend upon the network parameters. We already have some knowledge of the behavior of these trajectories from the previous analysis. Consistent with the mean field stability result, we expect that for an infinite oscillator system that $\varphi(t) = \Omega$. In other words, an isolated neuron with a predetermined phase and driving frequency will evolve as if ‘decoupled’ from the system and will not diffuse into the bath of neurons in the network. For a finite-size network, the deterministic, but disordered dynamics should produce an effective dynamical noise which is responsible for the population stability. We

can use the probabilistic description of the population to construct an effective stochastic Langevin equation for the underlying elements.

Suppose we isolate a ‘measured’ neuron

$$\dot{\phi}(t) = \Omega + \frac{K}{N} \sum_j \sin(\theta_j - \theta_i)$$

within a system of ‘unmeasured’ neurons

$$\dot{\theta}_i(t) = \omega_i + \frac{K}{N} \sum_j \sin(\theta_j - \theta_i) + \frac{K}{N} \sin(\phi(t) - \theta_i(t)).$$

One can write the action for the full system and then integrate over the other neurons, which we did in [35]. One can formally compute an effective action, which corresponds to the distribution of single neuron trajectories imposed by the distribution over initial conditions and driving frequencies. This action can be written equivalently as the following SDE

$$\dot{\phi}(t) = \Omega + \delta\Omega(t) + \xi(t) \quad (34)$$

where $\delta\Omega(t)$ is a frequency shift of the measured neuron due to interaction with the unmeasured neurons and $\xi(t)$ is a zero-mean stochastic forcing term, which will be in general non-Gaussian, non-white, and multiplicative (the amplitude depends on $\phi(t)$). In the stationary limit, the terms $\delta\Omega$ and ξ to $O(1/N)$ are given by

$$\delta\Omega_\infty = -\frac{1}{N} \frac{K^2}{2} \left[\frac{\Omega}{\Omega^2 + (\gamma - K/2)^2} \right] \left(\frac{4\gamma - K}{2\gamma - K} \right) \quad (35)$$

$$\langle \xi(t) \xi(t') \rangle = \frac{K^2}{2N} \left(\frac{2\gamma}{2\gamma - K} \right) e^{-(\gamma - K/2)(t - t')} \cos(\Omega(t - t')) \quad (36)$$

where we have assumed $t > t'$. The full calculation is given in [35]. The most important feature of this result is that the terms in the effective Langevin equation are given in terms of quantities which define the dynamics of the network. The noise term is not, for example, some externally supplied dynamical noise, but is given by a process which *mimics precisely* the ensemble of systems given by the distribution over network parameters and initial conditions.

We began section 3 with a discussion of Wilson–Cowan type ‘activity equations’ and how they could be extended to include fluctuation effects. Our Kuramoto example demonstrates the computation of fluctuation effects upon population properties and individual elements, but there is no obvious definition of ‘activity’ that might correspond to the Wilson–Cowan

activity variable. Neither the density of firing neurons nor the firing rate can serve as activity variables because the population dynamics do not depend directly on these quantities. Hence, closed activity equations cannot be written down for these variables. The closest parallel to an activity variable for phase coupled oscillators would be the population order parameter $Z(t)$, but this is not analogous to the quantity imagined in neuroscience, which is usually thought of as a ‘rate’ of neural firing or, perhaps, the local field potential.

A system of phase coupled neurons is not a natural microscopic system for an activity formulation because the interactions between the neurons are distributed across the entire firing cycle and thus no phase is particularly distinguished by any other. This phase invariance is broken in a synaptically coupled system, where neurons influence other neurons when they fire. Consider a system that interacts through a global synaptic variable:

$$\begin{aligned}\dot{\theta}_i &= f_i(\theta) + \alpha_i u(t) \\ \dot{u} + \beta u &= \frac{\beta}{N} \sum_j \delta(t - t_j^s)\end{aligned}$$

where the t_j^s represents the s th firing time of the j th neuron. If we consider ‘firing’ to be at $\theta \equiv \pi$, then $\theta_i(t_j^s) = \pi$. For the simple case where $f_i = I(t)$, the population density satisfies

$$\partial_t \eta + \partial_\theta [(I(t) + \alpha u(t))\eta] = 0 \quad (37)$$

$$\dot{u} + \beta u = \beta (I + \alpha u) \eta(\pi, t). \quad (38)$$

The firing rate is $\nu(t) = (I + \alpha u) \eta(\pi, t)$. We derived an action for this model in [36]. However, in this simple case, we can solve for the global activity variable in terms of the initial distribution of the phases. For concreteness, let us consider I to be a constant which is drawn from a distribution $g(I)$. Mean field theory is obtained by ignoring higher order moments (e.g. $\langle u\rho \rangle$). It has the same form as system (37) and (38) with u and η replaced by $\bar{u} = \langle u \rangle$ and $\rho = \langle \eta \rangle$, respectively. For, this simple case, the conservation equation can be solved in closed form so that the mean field theory for the global coupling can be written as

$$\dot{\bar{u}}(t) + \beta \bar{u}(t) = \beta \int d\Omega d\alpha (I_\Omega(t) + \alpha \bar{u}(t)) \rho_0 \left(\pi - \int_{t_0}^t dt' [I_\Omega(t') + \alpha \bar{u}(t')], \alpha, \Omega \right) \quad (39)$$

where ρ_0 is the initial condition of ρ . If ρ_0 is chosen to be a constant, then (39) would have the form of the Wilson–Cowan equation with a linear gain function and $\bar{u}$ acting as the activity variable. However, for general ρ_0 , $\bar{u}$ would depend on the entire history of the network dynamics.

It is straightforward to construct moments of the firing rate and global coupling from these expressions in the manner we used to derive the hierarchy. For this simple model, the firing rate fluctuations can be expressed as

$$\langle (\nu(t) - \bar{\nu}) (\nu(t) - \bar{\nu}) \rangle (\Delta t)^2 = -\frac{1}{N} \bar{\nu}^2 (\Delta t)^2 + \frac{1}{N} \bar{\nu} \Delta t (\delta(0) \Delta t). \quad (40)$$

The first term comes from the normalization for the unbiased estimator for the variance, whereas the second is a ‘Poisson’ term. For this simple model, the only fluctuation corrections are those which arise from standard finite sampling effects. In [36], we demonstrated that this computation can be extended to arbitrary $f(\theta)$, in particular we use the ‘theta model’ [26], with $f(\theta) = 1 - \cos \theta + (I + cu)(1 + \cos \theta)$. The phase dependence of θ has an important dependence upon the firing rate fluctuations. This produces an additional contribution to the firing rate fluctuations which cannot be expressed in closed form, except in terms of linear response functions of the population.

A natural question to ask is what happens when one includes heterogeneous connections into the network (i.e. quenched disorder). There are two answers. The first is that we can consider local populations of globally coupled oscillators and proceed as before. A Klimontovich equation and moment hierarchy exists for each population. If the individual populations are large then the hierarchy can be truncated with a system size expansion. However, even for small networks (even as small as one neuron), truncation may be possible using other expansion parameters such as the number of connections or a slow time scale. The second answer is that one can define a distribution in which the heterogeneous connections are drawn and average over the heterogeneity. If the network is self-averaging then this will still produce a consistent mean field theory and corrections to mean field theory. However, if it is not self-averaging, as in a spin glass, then such a mean field theory is not well defined as a description of a typical system. The ensemble distribution of systems will not be strongly peaked around a particular solution (e.g. the one defined by mean field theory). Owing to this, a system in such a regime may require a different approach. If some observables possessed well defined disorder averages then a replica averaging scheme could possibly be constructed. Most importantly, knowing when a system is self-averaging or not must be solved.

5. Discussion

We have shown how to go beyond mean field theory and include the effects of fluctuations and correlations by taking two approaches. The stochastic approach is more straightforward and leads directly to a set of self-consistent generalized activity equations for the mean rate and higher moments. The deterministic approach is better anchored to the underlying microscopic neuronal dynamics but requires a more circuitous route to a probabilistic description in which higher order statistics can be computed. One major computational difficulty in the deterministic approach is that mean field theory requires the solution of a nonlinear partial differential equation. For some simple models a closed form solution may be possible, but in general one will not be available. The higher order statistics are calculated by perturbing around the mean field solutions. Thus, a set of self-consistent generalized activity equations arising from a set of deterministic microscopic neuronal dynamics hinges on calculating mean field theory accurately. Even if a closed form solution of the mean field equation may not be available, the mean field system is easier to compute numerically than the original system, given that solving the mean field equations is

equivalent to ensemble averaging the original system over the entire distribution of possible initial conditions and parameter sets.

One technical hurdle to constructing closed generalized activity equations for the deterministic neurons is that the network often has perfect memory, in the same sense that the marginally stable infinite-size Kuramoto model has perfect memory. Equation (39) shows that Wilson–Cowan-like dynamics ensues if the stationary distribution of the network is flat or incoherent. Incoherence could be achieved by including stochastic effects into the neuronal dynamics. In this case memory is lost trivially by averaging over the noise distribution along with the initial data and parameters. A more intriguing hypothesis is that stability is induced through network heterogeneity. The ‘finite-size’ effects which we have demonstrated here are the expression of dynamic noise arising from network heterogeneity. If all oscillators in the Kuramoto model have the same driving frequency, i.e. $g(\omega) = \delta(\omega - \omega_0)$, then the network will always synchronize. It is the disorder that disrupts synchrony. In the mean field case, network coupling renders smooth perturbations stable in that the population order parameter will relax. Finite-size effects (i.e. heterogeneity) are necessary to stabilize marginal modes associated with the individual neuron spectrum. The implication for the pulse coupled networks is that heterogeneity induces a transition between an activity equation that displays the network ‘memory’ and an activity equation in which disorder induces a Wilson–Cowan-like behavior. We thus hypothesize that neural heterogeneity is necessary for ‘rate equations’ to be valid descriptions of population activity in neural systems.

There have been previous attempts to go beyond mean field theory. Brunel and Hakim [37] considered finite-size effects in a network where the connections were sparse enough so that the arrival times of synaptic events at a given neuron could be assumed to be uncorrelated. They used the ansatz that the inputs could be modeled by a Poisson process that was scaled by the number of inputs. We find that even for a fully connected network, the Poisson ansatz is partially correct to order $1/N$. Large deviation theory provides an explanation called ‘propagation of chaos’, where the uncertainty in the initial conditions is propagated forward by the deterministic dynamics of the system [38]–[41].

The mean field theory for our the deterministic microscopic system is comparable to the spike response theory [42, 43]. The use of phase oscillators allows for a continuity equation without a jump condition at the boundaries in a threshold crossing integrate-and-fire neuron. It may be possible to go beyond mean field within the spike response theory by incorporating the boundary conditions. Correlated neural activity due to finite-size effects have been explored by [44]–[46]. In [46], the authors develop a moment hierarchy for a Markov model of asynchronous irregular states of neural networks which is truncated through a combination of finite size and a scaling condition. Various kinetic theory and density approaches have been used to create moment equations to simplify large stochastic networks [47]–[50]. A mean field density Fokker–Planck approach to analyze properties and stability of the incoherent state of a stochastic network of neurons has been explored by numerous authors [33, 37], [51]–[58]. Golomb and Hansel [28] used this approach to study synchrony in sparse networks via a reduction of neuron dynamics to a phase model. Faugeras *et al* [59] constructed a mean field model using random weights and stochastic

inputs, but did not consider finite-size effects. Van Vreeswijk and Sompolinsky [60]–[62] demonstrated that a balanced network with sparse disordered connectivity can give rise to highly variable firing rates without stochastic input. They showed that, in the mean field limit, the fixed point of the network is a state where the mean excitatory and inhibitory contributions balance so that the fluctuations dominate the dynamics.

Acknowledgments

This research was supported by the Intramural Research Program of NIH/NIDDK.

References

1. Wilson H, Cowan J. *Biophys J*. 1972; 12:1. [PubMed: 4332108]
2. Wilson H, Cowan J. *Biol Cybernet*. 1973; 13:55.
3. Amari S. *Biol Cybernet*. 1977; 27:77.
4. Hopfield JJ. *Proc Nat Acad Sci*. 1982; 79:2552.
5. Cohen MA, Grossberg S. *IEEE Trans Syst Man Cybernet*. 1983; 13:815.
6. Coombes S. *Biol Cybernet*. 2005; 93:91.
7. Ermentrout GB, Cowan JD. *Biol Cybernet*. 1979; 34:137.
8. Bressloff PC, Cowan JD, Golubitsky M, Thomas PJ, Wiener MC. *Phil Trans R Soc B*. 2001; 356:299. [PubMed: 11316482]
9. Hopfield J. *Proc Nat Acad Sci*. 1984; 81:3088. [PubMed: 6587342]
10. Ermentrout B. *Rep Prog Phys*. 1999; 61:353.
11. Gerstner W. *Neural Comput*. 2000; 12:43. [PubMed: 10636933]
12. Gray C, Singer W. *Proc Nat Acad Sci*. 1989; 86:1698. [PubMed: 2922407]
13. Kopell N, Kramer M, Malerba P, Whittington M. *Front Hum Neurosci*. 2010; 4:187. [PubMed: 21103019]
14. Buice M, Cowan J. *Phys Rev E*. 2007; 75:051919.
15. Buice MA, Cowan JD, Chow CC. *Neural Comput*. 2010; 22:377. [PubMed: 19852585]
16. Bressloff PC. *SIAM J Appl Math*. 2010; 70:1488.
17. Doi M. *J Phys A: Math Gen*. 1976; 9:1465.
18. Doi M. *J Phys A: Math Gen*. 1976; 9:1479.
19. Peliti L. *J Physique*. 1985; 46:1469.
20. Janssen H, Täuber U. *Ann Phys*. 2005; 315:147.
21. Cornwall J, Jackiw R, Tomboulis E. *Phys Rev D*. 1974; 10:2428.
22. Cowan JD. *NIPS*. 1990:62–9.
23. Drossel B, Schwabl F. *Phys Rev Lett*. 1992; 69:1629. [PubMed: 10046273]
24. Buice M, Cowan J. *Prog Biophys Mol Biol*. 2009; 99:53. [PubMed: 19695282]
25. Albano, EV. Critical Behavior of a forest fire model with immune trees. 1994. <http://iopscience.iop.org/0305-4470/27/23/003>
26. Ermentrout G, Kopell N. *SIAM J Appl Math*. 1986:233–53.
27. Izhikevich, EM. *Dynamical Systems in Neuroscience, the Geometry of Excitability and Bursting*. Cambridge, MA: The MIT Press; 2007.
28. Golomb D, Hansel D. *Neural Comput*. 2000; 12:1095. [PubMed: 10905810]
29. Hildebrand EJ, Buice MA, Chow CC. *Phys Rev Lett*. 2007; 98:054101. [PubMed: 17358861]
30. Buice MA, Chow CC. *Phys Rev E*. 2007; 76:031118.
31. Martin P, Siggia E, Rose H. *Phys Rev A*. 1973; 8 :423.
32. Zinn-Justin, J. *Quantum Field Theory and Critical Phenomena*. Oxford: Oxford University Press; 2002.
33. Strogatz S, Mirollo R. *J Stat Phys*. 1991; 63:613.

34. Strogatz S, Mirollo R, Matthews P. Phys Rev Lett. 1992; 68:2730. [PubMed: 10045478]
35. Buice M, Chow C. Phys Rev E. 2011; 84:051120.
36. Buice MA, Chow CC. 2012 submitted.
37. Brunel N, Hakim V. Neural Comput. 1999; 11:1621. [PubMed: 10490941]
38. Gutkin E, Kac M. SIAM J Appl Math. 1983; 43:971.
39. McKean H Jr. Proc Nat Acad Sci. 1966; 56:1907. [PubMed: 16591437]
40. Baladron J, Fasoli D, Faugeras O, Touboul J. J Math Neurosci. 2012; 2:10. [PubMed: 22657695]
41. Touboul J. 2011 arXiv:1108.2414.
42. Gerstner W. Phys Rev E. 1995; 51:738.
43. Gerstner, W.; Kistler, W. Spiking Neuron Models: Single Neurons, Populations, Plasticity. Cambridge: Cambridge University Press; 2002.
44. Mattia M, Del Giudice P. Phys Rev E. 2002; 66:051917.
45. Soula H, Chow C. Neural Comput. 2007; 19:3262. [PubMed: 17970653]
46. El Boustani S, Destexhe A. Neural Comput. 2009; 21:46. [PubMed: 19210171]
47. Nykamp D, Tranchina D. J Comput Neurosci. 2000; 8:19. [PubMed: 10798498]
48. Nykamp D, Tranchina D. Neural Comput. 2001; 13:511. [PubMed: 11244554]
49. Cai D, Tao L, Shelley M, McLaughlin D. Proc Nat Acad Sci. 2004; 101:7757. [PubMed: 15131268]
50. Ly C, Tranchina D. Neural Comput. 2007; 19:2032. [PubMed: 17571938]
51. Abbott L, van Vreeswijk C. Phys Rev E. 1993; 48:1483.
52. Desai R, Zwanzig R. J Stat Phys. 1978; 19:1.
53. Mirollo R, Strogatz S. J Nonlinear Sci. 2007; 17:309.
54. Mirollo R, Strogatz S. Physica D. 2005; 205:249.
55. Treves A. Network. 1993; 4:259.
56. Fusi S, Mattia M. Neural Comput. 1999; 11:633. [PubMed: 10085424]
57. Fourcaud N, Brunel N. Neural Comput. 2002; 14:2057. [PubMed: 12184844]
58. Brunel N. J Comput Neurosci. 2000; 8:183. [PubMed: 10809012]
59. Faugeras O. Front Comput Neurosci. 2009; 3:1. [PubMed: 19255631]
60. Vreeswijk C, Sompolinsky H. Neural Comput. 1998; 10:1321. [PubMed: 9698348]
61. Vreeswijk C, Sompolinsky H. Sci New Series. 1996; 274:1724.
62. Sompolinsky H, Crisanti A, Sommers H. Phys Rev Lett. 1988; 61:259. [PubMed: 10039285]

Figure 1.

The effective spike count model. The figure shows a transition for the network, during which region I transitions from 5 to 4 active neurons, region II retains the same number of neurons, and region III transitions from 5 to 6 active neurons.

Figure 2.

The 2- and 3-state Cowan models. Each diagram shows the state transitions available to neurons in the network. Black arrows denote constant rates, while white arrows denote that the transitions are dependent upon the input to that neuron. The input is given by the number of active neurons in the network and the weight matrix, w_{ij} .

Figure 3.

Diagrammatic (Feynman) rules for the fluctuations about the mean. Time moves from right to left, as indicated by the arrow. The bare propagator $P_0(x, t|x', t')$ connects points at x' to x , where $x \equiv \{\theta, \omega\}$. Each branch of a vertex is labeled by x and x' and is connected to a factor of the propagator at x or x' . Each vertex represents an operator given to the right of that vertex. The '...' on which the derivatives act only include the incoming propagators, but not the outgoing ones. There are integrations over $\theta, \theta', \omega, \omega'$ and t at each vertex.

a)

b)

Figure 4.

Diagrams for the connected two-point function at tree level (a) and to one loop (b).