

Reduced models for binocular rivalry

Carlo R. Laing · Thomas Frewen ·
Ioannis G. Kevrekidis

Received: 27 August 2009 / Revised: 21 January 2010 / Accepted: 15 February 2010 / Published online: 25 February 2010
© Springer Science+Business Media, LLC 2010

Abstract Binocular rivalry occurs when two very different images are presented to the two eyes, but a subject perceives only one image at a given time. A number of computational models for binocular rivalry have been proposed; most can be categorised as either “rate” models, containing a small number of variables, or as more biophysically-realistic “spiking neuron” models. However, a principled derivation of a reduced model from a spiking model is lacking. We present two such derivations, one heuristic and a second using recently-developed data-mining techniques to extract a small number of “macroscopic” variables from the results of a spiking neuron model simulation. We also consider bifurcations that can occur as parameters are varied, and the role of noise in such systems. Our methods are applicable to a number of other models of interest.

Keywords Binocular rivalry · Diffusion map · Macroscopic · Data-mining

Action Editor: Carson C. Chow

C. R. Laing (✉)
IIMS, Massey University, Private Bag 102-904,
NSMC, Auckland, New Zealand
e-mail: c.r.laing@massey.ac.nz

T. Frewen
Department of Chemical Engineering, Princeton University,
Engineering Quadrangle, Olden Street, Princeton,
NJ 08544, USA

I. G. Kevrekidis
Department of Chemical Engineering, Princeton University,
PACM and Mathematics, Engineering Quadrangle,
Olden Street, Princeton, NJ 08544, USA

1 Introduction

Binocular rivalry occurs when two very different images are presented to the two eyes (Blake 2001; Blake and Logothetis 2002; Leopold and Logothetis 1999; Tong et al. 2006). Instead of perceiving the sum or average of the two images, the subject typically perceives only one image at a given time. However, there is normally repeated alternation between the two images (on the order of once every few seconds). The alternation is not exactly periodic, and the variability leads to a distribution of “dominance durations” (the times for which a particular image is perceived, before the other image is perceived). This phenomenon has been studied experimentally and computationally modelled for many years (Freeman 2005; Laing and Chow 2002; Moreno-Bote et al. 2007; Wilson 2003; Blake 2001; Grossberg et al. 2008). Some of these models involve a small number of variables (interpreted as “spiking rate” or “neural activity” of a subpopulation of neurons), and these models can be regarded as having been designed in such a way as to show the behaviour expected of them, i.e. slow oscillations (Dayan 1998; Freeman 2005; Stollenwerk and Bode 2003; Ashwin and Lavric 2010). Often, noise is added to these models to produce the observed distribution of dominance durations (Kalarickal and Marshall 2000; Lago-Fernandez and Deco 2002). Others have studied more realistic models of biophysically-based spiking neurons, with some of these authors also discussing rate models (Laing and Chow 2002; Moreno-Bote et al. 2007; Wilson 2003).

However, a systematic derivation of a rate model for binocular rivalry from a spiking model is still lacking. Indeed, the derivation of accurate “macroscopic”

models from detailed “microscopic” models remains one of the outstanding problems in computational neuroscience; for various approaches see, for example, Wilson and Cowan (1972), Ch. 6 in Gerstner and Kistler (2002), Cai et al. (2004) and Tranchina (2009). We perform two such derivations here, using intuition in the selection of macroscopic variables and the approach described in Gradišek et al. (2000) for the analysis of stochastic systems. We first use our experience to choose several macroscopic variables, and process the results of a long simulation to extract deterministic and stochastic components of equations governing the dynamics of these variables. We then demonstrate that similar results can be obtained by blindly “data-mining” the results of a long simulation, to automatically extract appropriate macroscopic variables. We use the recently-developed diffusion map approach (Coifman and Lafon 2006; Erban et al. 2007; Nadler et al. 2006) to accomplish this.

Note that unlike, for example, Cai et al. (2004) and Tranchina (2009), we do not provide an analytical reduction from a microscopic model to a macroscopic one. Instead, our reduction is numerical and requires the simulation of the detailed microscopic model at the parameter value(s) of interest. This paper concerns the processing of the results of such simulations to obtain, among other things, estimates of the functions used in a macroscopic model.

In Section 4 we investigate the effects of varying parameters in the system and show that we can move from an “oscillator model” to an “attractor model” (Moreno-Bote et al. 2007) by changing a single parameter. In Section 5 we add noise to the system and show that this affects both the deterministic component and the full stochastic dynamics of the macroscopic model that we derive. Section 6 shows how to initialise the system consistent with specific values of the macroscopic variables, and we conclude in Section 7.

2 Fine-scale model

The model we use is a slight modification of that previously described by Laing and Chow (2002). We refer to it as the fine-scale model and briefly describe it here. The model consists of two populations of Hodgkin–Huxley-like neurons, excitatory and inhibitory. Each population has 60 neurons, each of which has a “preferred orientation,” so that it fires at its maximal rate when either eye is presented with a grating of that orientation. The preferred orientations are chosen uniformly from the range $[0, 180^\circ]$, and thus the neurons in each population can be thought of as lying on a

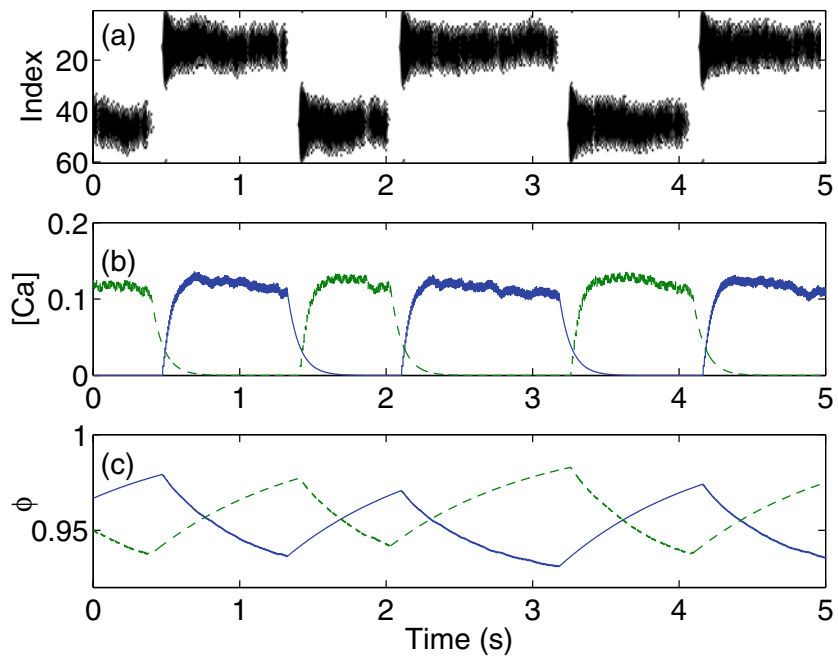
ring, with position around the ring corresponding to preferred orientation (see Fig. 1 of Laing and Chow 2002). We propose that a given percept is represented as a localised group of active neurons (Gutkin et al. 2001; Laing and Chow 2001). We model the situation in which the two eyes are presented with orthogonal gratings (Logothetis et al. 1996), and thus inject current at two locations in the network centred around neurons whose preferred orientation differs by 90° , i.e. on opposite sides of the ring.

We also assume that excitatory neurons are synaptically coupled with a strength which is a Gaussian function of the difference between their preferred orientations. There is a similar type of coupling from excitatory neurons to inhibitory neurons, within the inhibitory population, and from inhibitory to excitatory, with coupling strength always being a Gaussian function of the difference between preferred orientations. See Appendix for details.

We include two slow processes in our fine-scale model. The first is spike frequency adaptation in the excitatory neurons; this has a time-constant of 80 ms. The second is synaptic depression in the excitatory–excitatory connections, with a time-constant of 1 s. Because of the form of the injected current (see Appendix) the neurons labelled 1–30 are associated with one percept and those labelled 31–60 are associated with the other. The network robustly undergoes oscillations between one half of the network being active and the other half being active, switching every second or two. See Fig. 1 for an example, or Laing and Chow (2002). The changing calcium concentration is responsible for spike frequency adaptation, while ϕ is the variable controlling synaptic depression. Laing and Chow (2002) showed that their model reproduced a number of experimentally observed phenomena.

While this fine-scale model is detailed and biophysically realistic (and could be made more detailed and realistic) it is clear that from a macroscopic point of view, the network is undergoing noisy oscillations. However, it is not clear whether the system should be thought of as an intrinsic oscillator with added noise, or as a bistable system, which alternates between states purely as a result of being driven by noise (Moreno-Bote et al. 2007; Shpiro et al. 2009). It is this observation of noisy oscillations at the macroscopic level that is behind the study of various rate models for rivalry which have few variables (Kalarickal and Marshall 2000; Lago-Fernandez and Deco 2002; Laing and Chow 2002; Ashwin and Lavric 2010). However, a systematic derivation of a rate model from a detailed spiking model has not previously been performed, and it is this question that we now address.

Fig. 1 (a): A rasterplot for the excitatory population. Each firing of an excitatory neuron is marked by a *black bar*. (b): [Ca] for neurons 15 (*solid*) and 45 (*dashed*). (c): ϕ for neurons 15 (*solid*) and 45 (*dashed*). The activity in the inhibitory population mimics that in the excitatory



3 Macroscopic models

3.1 Deriving a macroscopic model, choosing macroscopic variables

Based on our observations of the network, i.e. alternation of activity in the two halves of the population, we define two macroscopic variables: χ is defined to be the mean of [Ca] for neurons 31–60, minus the mean of [Ca] for neurons 1–30. Φ is defined to be the mean of ϕ for neurons 31–60, minus the mean of ϕ for neurons 1–30. We thus expect χ to be a fast variable, relative to Φ . Typical behaviour of χ and Φ is shown in Fig. 2. We see that in terms of these variables, the system possesses a stable, “noisy limit cycle,” and it seems reasonable that another system which has a similar stable, noisy limit cycle could be an appropriate model for the microscopic model under investigation.

We do not attempt to analytically derive deterministic nor stochastic differential equations governing the dynamics of Φ and χ from fine-scale Eqs. (31)–(40) given in Appendix. Instead we take the “equation-free” approach of assuming that such equations exist and then estimating the right-hand sides of those equations (Kevrekidis et al. 2003; Laing 2006). We do this by processing data like that shown in Fig. 2 to extract estimates of the terms involved in stochastic DEs (SDEs) for χ and Φ using the techniques in Gradišek et al. (2000) (see also Friedrich et al. 2000; Laing et al. 2007; van Mourik et al. 2006). These SDEs will form our macroscopic model, and are assumed to linearly

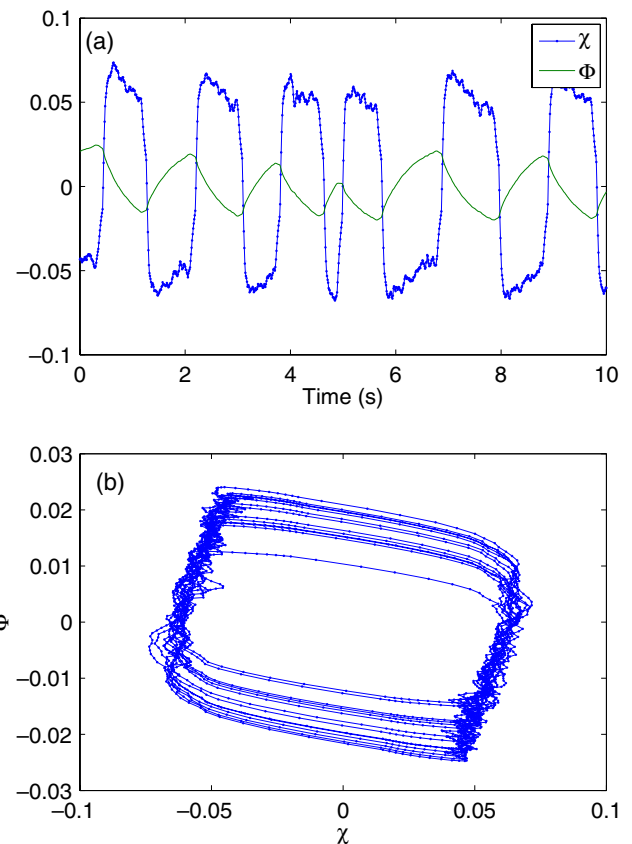


Fig. 2 (a): Φ and χ as functions of time, directly from a simulation of the fine-scale model. (b): Motion in the Φ, χ plane for a simulation of length 30 s. Motion is in the clockwise direction

combine purely deterministic and purely stochastic components, i.e. they form a vector Langevin equation. It should be noted that the fine-scale (Hodgkin–Huxley-like) system is purely deterministic, with apparently stochastic switching resulting from finite-size effects (non-synchronous, individual synaptic inputs) and the highly nonlinear nature of the full system. Also, it is an *assumption* that such SDEs provide an accurate description of the system. In general, the validity of this assumption can be checked by comparing results from these SDEs with those from the full, detailed system (Eqs. (31)–(40))—see the end of this section for further discussion.

We define

$$\mathbf{X} = \begin{pmatrix} \chi \\ \Phi \end{pmatrix} \in \mathbb{R}^2 \tag{1}$$

and assume that \mathbf{X} satisfies an as yet unknown Langevin equation. It follows that the probability density $P(\mathbf{X}, t)$ satisfies a Fokker–Planck equation:

$$\frac{\partial P(\mathbf{X}, t)}{\partial t} = \left[-\sum_i \frac{\partial}{\partial X_i} f_i(\mathbf{X}) + \frac{1}{2} \sum_{i,j} \frac{\partial^2}{\partial X_i \partial X_j} D_{ij}(\mathbf{X}) \right] P(\mathbf{X}, t). \tag{2}$$

The functions f_i and D_{ij} are defined from the dynamics through the following formulas:

$$f_i(\mathbf{Y}) = \lim_{\Delta t \rightarrow 0} \frac{\langle X_i(t + \Delta t) - X_i(t) \rangle}{\Delta t} \Big|_{\mathbf{X}(t)=\mathbf{Y}} \tag{3}$$

and

$$D_{ij}(\mathbf{Y}) = \lim_{\Delta t \rightarrow 0} \frac{\langle (X_i(t + \Delta t) - X_i(t))(X_j(t + \Delta t) - X_j(t)) \rangle}{\Delta t} \Big|_{\mathbf{X}(t)=\mathbf{Y}} \tag{4}$$

where the angled brackets denote expectation. Note from its definition that D is symmetric. The vector

$$\mathbf{f}(\mathbf{X}) = \begin{pmatrix} f_1(\chi, \Phi) \\ f_2(\chi, \Phi) \end{pmatrix} \in \mathbb{R}^2 \tag{5}$$

governs the deterministic part of the macroscopic dynamics, while $D(\mathbf{X})$ is related to the stochastic aspect of the Langevin equation for \mathbf{X} . We estimate $\mathbf{f}(\mathbf{X})$ and $D(\mathbf{X})$ from a finite-length simulation, using the following expressions with finite Δt (in practice, we used $\Delta t = 10$ ms).

$$f_i(\mathbf{Y}) \approx \frac{\langle X_i(t + \Delta t) - X_i(t) \rangle}{\Delta t} \Big|_{\mathbf{X}(t)=\mathbf{Y}} \tag{6}$$

and

$$D_{ij}(\mathbf{Y}) \approx \frac{\langle (X_i(t + \Delta t) - X_i(t))(X_j(t + \Delta t) - X_j(t)) \rangle}{\Delta t} \Big|_{\mathbf{X}(t)=\mathbf{Y}} - \Delta t f_i(\mathbf{Y}) f_j(\mathbf{Y}) \tag{7}$$

where the last term in Eq. (7) helps to correct for the finite size of Δt (Ragwitz and Kantz 2001). Having a finite amount of data we also have to partition the relevant part of the phase plane into equal-sized rectangles, the result being that we obtain estimates for $\mathbf{f}(\mathbf{X})$ and $D(\mathbf{X})$ at a finite number of points in the \mathbf{X} plane (Kuusela et al. 2003). Figure 3 shows f_1 and f_2 extracted from a simulation of Eqs. (31)–(40) of duration 500 s. Note that the maximum magnitude of f_1 is more than ten times that of f_2 . Figure 4 shows




Fig. 3 Estimates of (a): $f_1(\mathbf{X})$ and (b): $f_2(\mathbf{X})$, i.e. the components of the vector in Eq. (5)




Fig. 4 Plots of (a) f_1 and (b) f_2 as function of χ for the two specified values of Φ (-0.015 and 0.015)

horizontal “slices” through the images shown in Fig. 3 at two particular value of Φ .

Since D is symmetric we represent it by its Cholesky decomposition: $D = GG^T$, where

$$G = \begin{pmatrix} G_{11} & 0 \\ G_{21} & G_{22} \end{pmatrix}. \tag{8}$$

We have

$$G_{11} = \sqrt{D_{11}} \tag{9}$$

$$G_{21} = D_{12}/G_{11} \tag{10}$$

$$G_{22} = \sqrt{D_{22} - G_{21}^2} \tag{11}$$

Figure 5 shows estimates of $G_{11}(\mathbf{X})$, $G_{21}(\mathbf{X})$ and $G_{22}(\mathbf{X})$, extracted as above. Note that the maximum magnitude of G_{11} is approximately 10 times the maximum magnitude of G_{21} , which is approximately 10 times the maximum magnitude of G_{22} . The grid used

in Figs. 3 and 5 has 19 equally spaced points over the range of χ values shown and 29 equally spaced points over the range of Φ values shown.




Fig. 5 Estimates of (a): $G_{11}(\mathbf{X})$, (b): $G_{21}(\mathbf{X})$ and (c): $G_{22}(\mathbf{X})$

Having determined $\mathbf{f}(\mathbf{X})$ and D , we can, for example, simulate the deterministic component of the macroscopic dynamics:

$$\frac{d\mathbf{X}}{dt} = \mathbf{f}(\mathbf{X}). \tag{12}$$

The results are shown in Fig. 6 (left column). We see that the deterministic dynamics is characterised by a globally attracting limit cycle surrounding an unstable fixed point.

We can also simulate the full Langevin equation for \mathbf{X} . Choosing a small time step Δt , the Euler–Maruyama scheme is

$$\begin{aligned} \mathbf{X}((i+1)\Delta t) \\ = \mathbf{X}(i\Delta t) + \Delta t \mathbf{f}(\mathbf{X}(i\Delta t)) + \sqrt{\Delta t} G(\mathbf{X}(i\Delta t)) \Gamma^i \end{aligned} \tag{13}$$

where

$$\Gamma^i = \begin{pmatrix} \Gamma_1^i \\ \Gamma_2^i \end{pmatrix} \in \mathbb{R}^2 \tag{14}$$

and Γ_1^i and Γ_2^i are independently selected from a normal distribution with mean zero and variance 1. For values of \mathbf{X} not on the grid mentioned above, linear interpolation of the values of $\mathbf{f}(\mathbf{X})$ and $G(\mathbf{X})$ at these grid points is used in the simulations above. In the course of the simulation of Eq. (13), \mathbf{X} may enter a region which the original simulation never reached, for which we do not have estimates of \mathbf{f} and G . In these regions we (somewhat arbitrarily) set $G = 0$ and $\mathbf{f} = -0.2\mathbf{X}$. The effect of this is to move the trajectory gradually towards the origin, so that it re-enters the region for which estimates of \mathbf{f} and G are available. The results of such a simulation are shown in Fig. 6 (right column). The agreement appears to be quite good. Needless to say, such simulations are much faster than simulations of the original system, with the generation of the random vectors Γ^i consuming most of the computational effort. The issue of finding estimates of \mathbf{f} and G for values of \mathbf{X} not reached by the original simulation is addressed in Section 6.

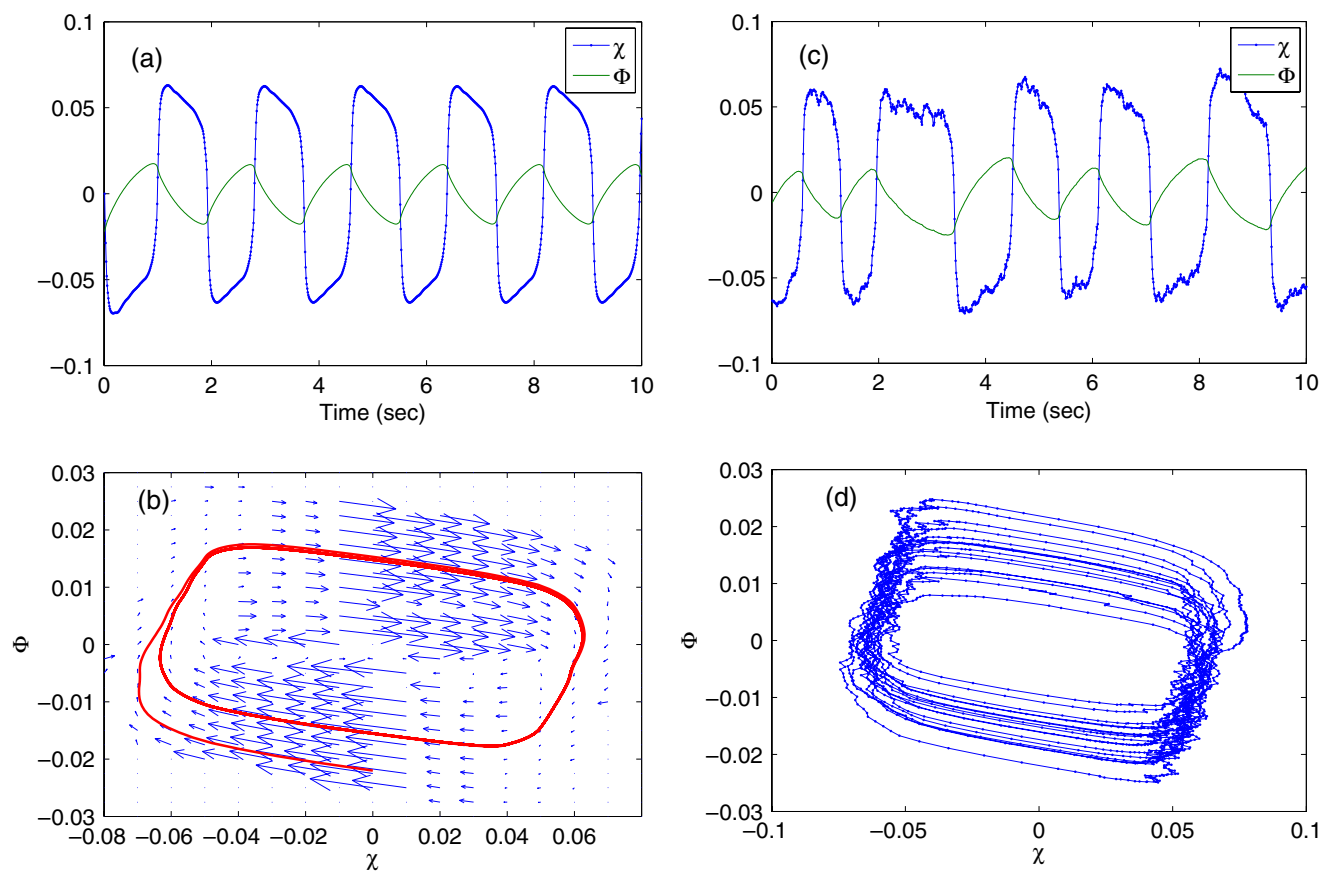


Fig. 6 Left column: a simulation of the deterministic dynamics (Eq. (12)), including transients. (a): Φ and χ as functions of time. (b): the trajectory in the top panel laid over the deterministic direction field. Right column: Simulations of the Langevin

equation (Eq. (13)). (c): Φ and χ as functions of time. (d): Motion in the Φ, χ plane for a simulation of length 30 s. Compare with Fig. 2

It is possible to check the validity of approximating the dynamics shown in Fig. 2 by a Langevin equation for \mathbf{X} , as we have done. We can choose a point $\hat{\mathbf{X}}$ in phase space and consider all points $\{\mathbf{X}(t_1), \mathbf{X}(t_2), \dots, \mathbf{X}(t_m)\}$ that come close to it during a finite simulation, and then examine the distribution of differences

$$\{\mathbf{X}(t_1) - \mathbf{X}(t_1 + \Delta t), \mathbf{X}(t_2) - \mathbf{X}(t_2 + \Delta t), \dots, \mathbf{X}(t_m) - \mathbf{X}(t_m + \Delta t)\}. \tag{15}$$

If, for all points $\hat{\mathbf{X}}$, the corresponding distribution of differences forms a bivariate normal distribution, a Langevin equation is appropriate. The parameters of the bivariate normal distribution are directly related to the values of \mathbf{f} and G (Kuusela et al. 2003).

As an example we chose $\hat{\mathbf{X}} = (\hat{\chi} \ \hat{\Phi})^T = (-0.06 \ 0.005)^T$ and examined all points for which $(\chi, \Phi) \in (-0.065, -0.055) \times (0.0025, 0.0075)$, using these points and the values of χ and Φ at a time Δt later to form the differences (Eq. (15)). Figure 7(a) shows the distribution of these differences. Figure 7(c) shows a histogram of the χ values for these differences, while panel (d) shows a histogram of the corresponding Φ values. Also shown in panel (c) is a normal distribution with mean $f_1(\hat{\chi}, \hat{\Phi})\Delta t$ and standard deviation $G_{11}(\hat{\chi}, \hat{\Phi})\sqrt{\Delta t}$, and in panel (d) a normal distribution with mean $f_2(\hat{\chi}, \hat{\Phi})\Delta t$ and standard deviation $\sqrt{[G_{21}(\hat{\chi}, \hat{\Phi})]^2\Delta t + [G_{22}(\hat{\chi}, \hat{\Phi})]^2\Delta t}$. We see that both


distributions are well-fit by normal distributions, with parameters corresponding to the estimated values of \mathbf{f} and G . The correlation coefficient between the values of χ and the values of Φ for the points in Fig. 7(a) is -0.96397 , in excellent agreement with the theoretical value of $G_{21}(\hat{\chi}, \hat{\Phi})/\sqrt{[G_{21}(\hat{\chi}, \hat{\Phi})]^2 + [G_{22}(\hat{\chi}, \hat{\Phi})]^2} = -0.96397$. Repeating this analysis for other values of $\hat{\mathbf{X}}$ gives similar results (not shown).

The Markovian nature of the dynamics can also be checked (Bahraminasab et al. 2008). Given the time-series for \mathbf{X} and our estimates of \mathbf{f} and G , one can, for example, solve Eq. (13) at each time-step for the corresponding values of Γ_1^i and Γ_2^i . If the dynamics are Markovian we expect these values to be δ -correlated in time. Figure 7(b) shows the autocorrelation of the sequence $\{\Gamma_1^i\}$. We see that there are some weak correlations beyond one time-step of $\Delta t = 10$ ms, but the rapid decay supports the use of a Langevin equation to model the dynamics of \mathbf{X} .

3.2 Deriving a macroscopic model using data mining

While the method above was very successful in suggesting a macroscopic model, it relied on us choosing appropriate variables with which to form the macroscopic model, namely χ and Φ . We were able to do this successfully because of our experience simulating

Fig. 7 (a): Distribution of differences (Eq. (15)). (b): Autocorrelation of the sequence $\{\Gamma_1^i\}$. (c) and (d) show χ and Φ values, respectively, of the data in panel (a), together with normal distributions with parameter values given by \mathbf{f} and G evaluated at $\hat{\mathbf{X}}$. See text for more details



the model, and also because we had access to the fine-scale model, both in terms of the governing equations, and being able to simulate them. However, in many situations this will not be the case. For example, our fine-scale simulator may be a “legacy code” which we are not able to inspect or modify, or which was not fully documented. Alternatively, the results of a simulation may be too complicated (e.g. having too many degrees of freedom) for us to use our intuition to correctly choose appropriate coarse-grained observables (variables). Instead, we would like an automated method which can determine, from the results of a fine-scale simulation, variables which can be used to describe the dynamics at a macroscopic level.

One approach is to “mine” the data collected from one or more long, detailed simulations in order to find an accurate low-dimensional description of it. We then use coordinates in this low-dimensional space as macroscopic variables. Since we have dynamics on the high-dimensional data set, this allows us to observe the dynamics of the macroscopic variables, thus providing, in effect, a low-dimensional model. The assumption behind this approach is that, in fact, there is such a low-dimensional description of the dynamics. It is similar in spirit to centre manifold calculations in the neighbourhood of a bifurcation, or to the use of approximate inertial manifolds (Jolly et al. 1990; Rega and Troger 2005). We will perform this data-mining (or “manifold learning”) using the recently-developed diffusion map approach (Coifman and Lafon 2006; Erban et al. 2007; Laing et al. 2007; Nadler et al. 2006).

3.2.1 Diffusion maps

In this section we briefly describe the application of diffusion maps to our problem. The procedure can be thought of as a nonlinear generalization of principal component analysis (Jolliffe 2002) and further details can be found in Erban et al. (2007) and Laing et al. (2007).

The behaviour of the inhibitory network will mimic that of the excitatory, so we do not consider the inhibitory neuron variables. We sample the 360 variables associated with the excitatory network ($V_e^k, n_e^k, h_e^k, s_e^k, [Ca]^k, \phi^k$) for $k = 1, \dots, 60$ at $N = 4,000$ equally spaced timesteps (10 ms apart). We thus have N vectors in \mathbb{R}^{360} . Because the variables have different scales (and indeed units) we first subtract the mean of each variable:

$$X^k(j) \rightarrow X^k(j) - \mu_X^k \quad (16)$$

for each $k \in \{1, \dots, 60\}$ and $X \in \{V_e, n_e, h_e, s_e, [Ca], \phi\}$ where

$$\mu_X^k = \frac{1}{N} \sum_{j=1}^N X^k(j) \quad (17)$$

and $X^k(j)$ is the value of X^k at the j th time point. We then rescale:

$$X^k(j) \rightarrow \frac{X^k(j)}{M_X^k - m_X^k} \quad (18)$$

where

$$M_X^k = \max_j X^k(j) \text{ and } m_X^k = \min_j X^k(j). \quad (19)$$

We refer to these shifted and rescaled vectors as $\{\mathbf{x}_i\}_{i=1, \dots, N}$. We then construct a similarity matrix K :

$$K_{i,j} = \exp \left[- \left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\epsilon} \right)^2 \right] \quad (20)$$

where $\|\cdot\|$ indicates the Euclidean norm and ϵ is a characteristic distance; here we use $\epsilon = \sqrt{50}$. We create a diagonal normalisation matrix D , where

$$D_{i,i} = \sum_{j=1}^N K_{i,j} \quad (21)$$

and the Markov matrix $M = D^{-1}K$. M has eigenvalues $1 = \lambda_0 \geq \lambda_1 \geq \dots \lambda_{N-1} \geq 0$ with right eigenvectors $\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{N-1}$. Note that the eigenvector \mathbf{v}_0 is constant (all of its entries are 1). If there is a spectral gap after a few eigenvalues, this suggests that the data is low-dimensional, and the components of the data points on the leading eigenvectors, $\mathbf{v}_1, \mathbf{v}_2, \dots$ provide a useful low-dimensional representation of the data set $\{\mathbf{x}_i\}$ (Coifman and Lafon 2006; Nadler et al. 2006). Here we use only the first two eigenvectors. The diffusion map is then

$$\mathbf{x}_i \mapsto \begin{pmatrix} \mathbf{v}_1^{(i)} \\ \mathbf{v}_2^{(i)} \end{pmatrix} \in \mathbb{R}^2 \quad (22)$$

for $i = 1, \dots, N$, where $\mathbf{v}_a^{(i)}$ is the i th component of \mathbf{v}_a . This is a mapping from \mathbb{R}^{360} to \mathbb{R}^2 . Note that this procedure does not use the dynamics of the \mathbf{x}_i . We refer to the unbolded v_1 and v_2 as “diffusion map coordinates,” and regard them as scalar variables. Thus, for example, $\mathbf{v}_1^{(i)}$ is a specific value of v_1 .

In Fig. 8 we plot v_1, v_2 and v_3 as functions of time for the 4,000 data points used to construct the diffusion map, together with χ and Φ for the same data points. The first coordinate v_1 clearly captures the behaviour of χ , while v_2 captures the dynamics of Φ . The meaning of v_3 is less clear, as it varies the same way, irrespective

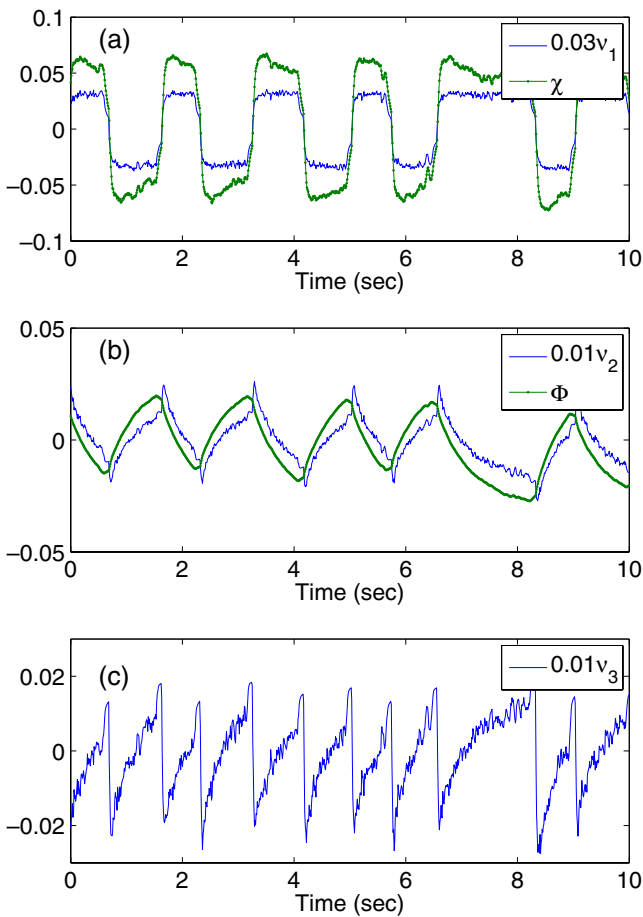


Fig. 8 (a): Diffusion map coordinate v_1 (multiplied by 0.03) and χ as functions of time. (b): Diffusion map coordinate v_2 (multiplied by 0.01) and Φ as functions of time. (c): Diffusion map coordinate v_3 (multiplied by 0.01)

of whether v_2 is increasing or decreasing. The variables v_1 and v_2 were extracted in an almost completely automated fashion, and their correspondence with χ and Φ can be regarded as confirmation that these were indeed good macroscopic variables.

Several points regarding the results in this section should be discussed. One point of interest is the sampling rate of 100 Hz. One might be concerned that such a rate is not fast enough to fully resolve individual action potentials, leading somehow to v_1 and v_2 only reflecting the dynamics of the slow variables. However, this is not the case. The behaviour of the network is driven by the states of the slow variables, and individual action potentials do not need to be resolved for the method to work. All we require is that we sample for long enough to cover several switches of activity between sides of the fine-scale network, and that samples be close enough in time that we can reliably estimate time derivatives of v_1 and v_2 using finite differences.

As a check we repeated the analysis in this section but sampling the fine-scale model at 1,000 Hz, and obtained essentially the same results (not shown). Another issue is the use of a nonlinear manifold learning algorithm as opposed to a linear method such as principal components analysis. For the data set used here, i.e. $\{\mathbf{x}_i\}_{i=1,\dots,N}$, the first two principal components are actually very similar to v_1 and v_2 (not shown). However, in the absence of any information regarding the structure of the data set to be analyzed, the most general algorithm (i.e. a nonlinear one) should clearly be used.

3.2.2 The Nyström formula

Now that we have identified several variables with which to describe the macroscopic dynamics of the system, we would like to do the same as before, i.e. estimate functions associated with the deterministic and stochastic dynamics of those variables. However, we only have the values of v_1 and v_2 for $N = 4,000$ data points. In practice, we cannot significantly increase this number, since the matrix K has dimensions $N \times N$. However, based on the assumption (which is supported by the computed spectrum of M) that the data are low-dimensional, the Nyström formula for eigenspace interpolation can be used to find the values of v_1 and v_2 for any number of other data points not used in the construction of K (Bengio et al. 2004; Erban et al. 2007). This gives us more v_1, v_2 pairs with which to estimate the deterministic and stochastic dynamics of v_1 and v_2 .

To do this, we note that $M = D^{-1/2}SD^{1/2}$, where $S = D^{-1/2}KD^{-1/2}$. M and S are thus similar, and therefore have the same eigenvalues, $\lambda_0 \geq \lambda_1 \geq \dots \lambda_{N-1}$. Let $\{\mathbf{U}_j\}_{j=0,\dots,N-1}$ be the corresponding eigenvectors of S . These are related to the eigenvectors of M through

$$\mathbf{v}_j = D^{-1/2}\mathbf{U}_j. \tag{23}$$

We have $S\mathbf{U}_j = \lambda_j\mathbf{U}_j$, or

$$\mathbf{U}_j^{(i)} = \frac{1}{\lambda_j} \sum_{k=1}^N S_{i,k} \mathbf{U}_j^{(k)} \tag{24}$$

where $\mathbf{U}_j^{(i)}$ is the i th component of \mathbf{U}_j . Suppose we have new vector \mathbf{x}_{new} , and we want to know the values of v_1 and v_2 associated with it. We create an $N \times 1$ vector \mathbf{K}_{new} whose k th component is

$$\mathbf{K}_{new}^{(k)} = \exp \left[- \left(\frac{\|\mathbf{x}_{new} - \mathbf{x}_k\|}{\epsilon} \right)^2 \right]. \tag{25}$$

We also have a generalised kernel vector \mathbf{S}_{new} whose i th entry is

$$\mathbf{S}_{new}^{(i)} = \left(\sum_{k=1}^N K_{i,k} \right)^{-1/2} \left(\sum_{k=1}^N \mathbf{K}_{new}^{(k)} \right)^{-1/2} \mathbf{K}_{new}^{(i)}. \quad (26)$$

The entries in \mathbf{S}_{new} quantify the pairwise similarities between \mathbf{x}_{new} and the vectors in $\{\mathbf{x}_k\}$, consistent with the definition of S . The eigenvector component \mathbf{U}_j^{new} corresponding to \mathbf{x}_{new} is then

$$\mathbf{U}_j^{new} = \frac{1}{\lambda_j} \sum_{i=1}^N \mathbf{S}_{new}^{(i)} \mathbf{U}_j^{(i)} \quad (27)$$




Fig. 9 Estimates of the functions governing the deterministic dynamics of v_1 and v_2 . (a): f_1 ; (b): f_2




Fig. 10 Estimates of (a): G_{11} , (b): G_{21} and (c): G_{22} as functions of the variables v_1 and v_2

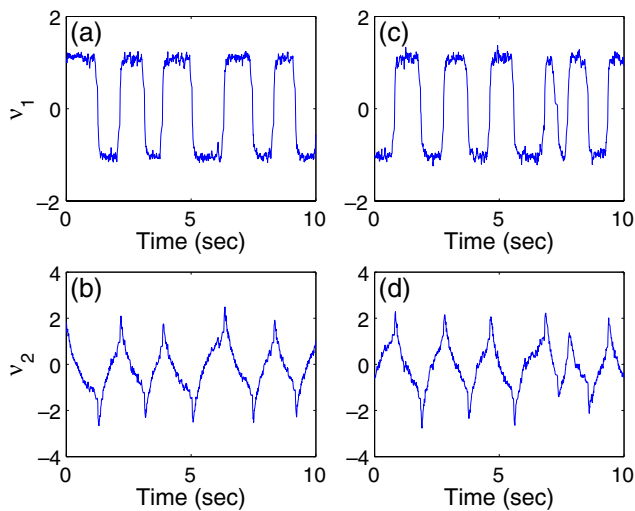


Fig. 11 *Left column:* (a) and (b) show v_1 and v_2 , respectively, as extracted from a simulation of the microscopic model. *Right column:* (c) and (d) show results of a Langevin simulation (i.e. a sample path) of v_1 and v_2 , respectively, using the functions shown in Figs 9 and 10

for $j = 0, 1, 2$ (since we only want the values of v_1 and v_2) and the value of v_j corresponding to \mathbf{x}_{new} is

$$v_j^{new} = \frac{\mathbf{U}_j^{new}}{\mathbf{U}_0^{new}} \quad (28)$$

Performing this procedure for data points not used in the construction of K , we obtain a large number of v_1, v_2 pairs from which we can estimate the vector-valued function $\mathbf{f}(v_1, v_2)$, governing the deterministic dynamics of v_1 and v_2 , and the matrix $D(v_1, v_2)$, corresponding to the stochastic component of the dynamics. Figure 9 shows estimates of the components of \mathbf{f} , and Fig. 10 shows estimates of the coefficients G_{11}, G_{21}, G_{22} (associated with the matrix D), estimated from a total of 50,000 data points. Once we have these functions, we can simulate the corresponding Langevin equation for v_1 and v_2 , as before—see Fig. 11. The comparison with dynamics of the fine-scale simulation is again satisfactory.

4 Changing parameters

So far we have shown how to derive a macroscopic model for the dynamics, using either “experience-based” coordinates, or those obtained by data-mining the results of a simulation, for fixed parameters. We are also interested in the effects of changing parameters in the model, as that will cause the dynamics to change. For example, if we decrease B (the parameter control-

ling the strength of synaptic depression) from 1.3 to 0.8, the average dominance duration increases, but we still obtain oscillations—see Fig. 12(a), (b). Using the variables Φ and χ defined as before, we can extract

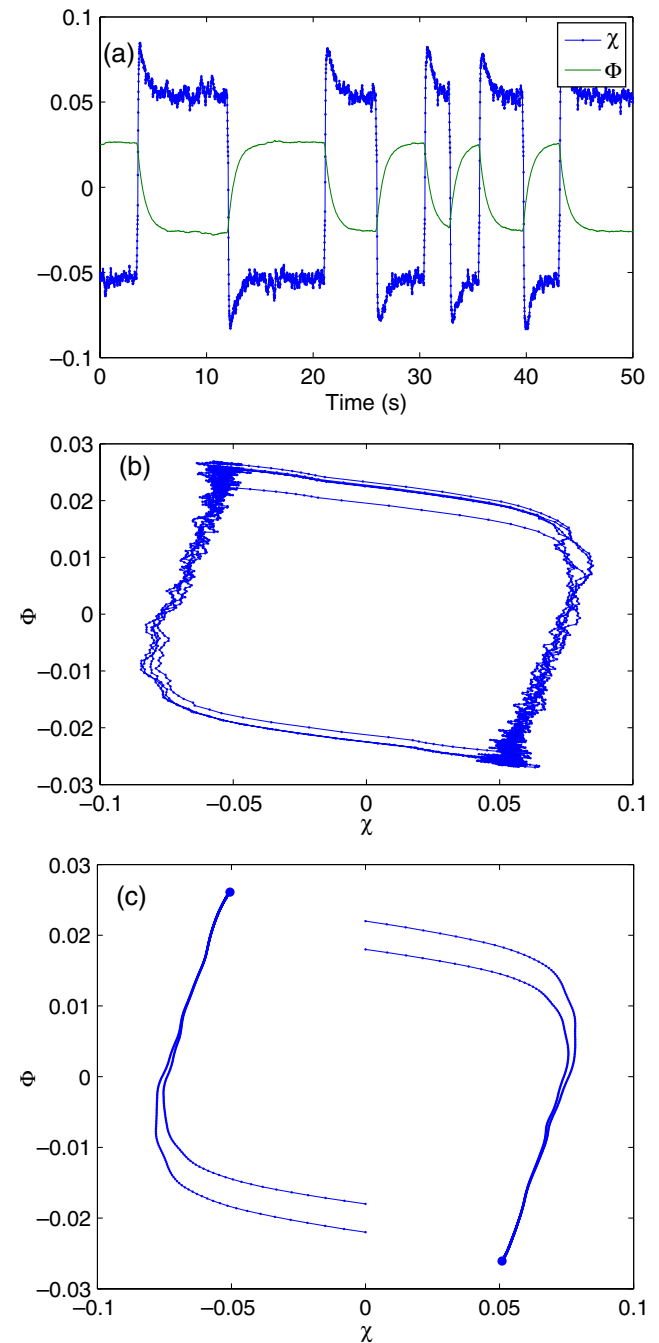


Fig. 12 (a) and (b): Dynamics of the fine-scale model for $B = 0.8$. (a): Φ and χ as functions of time. (b): motion in the χ, Φ plane. (c): Trajectories from several simulations of the deterministic system (Eq. (12)) when $B = 0.8$. The attractors are shown with large filled circles. Rotation is clockwise in panels (b) and (c)

estimates of \mathbf{f} and G_{11} , G_{21} , G_{22} , for this new parameter value. What is interesting in this case is the deterministic dynamics, as shown in Fig. 12(c). The deterministic dynamics no longer have a stable limit cycle. Instead, there are two symmetry-related stable fixed points. The deterministic component of the dynamics at the macroscopic scale has undergone a bifurcation, even though the overall stochastic dynamics does not appear to have done so.

By using linear interpolation in phase space on the components of \mathbf{f} as estimated at discrete points in the χ, Φ plane, we can estimate the location of the unstable fixed points of the deterministic component of the dynamics as well. The unstable fixed points are saddles, with one positive and one negative eigenvalue, as expected from their creation in a symmetric pair of saddle-node bifurcations as B is decreased. These bifurcations are similar to the saddle-node-on-an-invariant-circle (SNIC), or saddle-node-infinite-period (SNIPER) bifurcation, the difference being that the symmetry of the system forces two pairs of fixed points (related under the transformation $(\chi, \Phi) \mapsto (-\chi, -\Phi)$) to appear at the same parameter value.

Decreasing B still further would cause an apparent bifurcation in the stochastic dynamics, in the sense that rather than switching between two states the system would remain practically forever at one noisy fixed point (not shown).

The underlying cause of switching in binocular rivalry is as yet unknown. Moreno-Bote et al. (2007) recently studied a stochastic model for binocular rivalry that they referred to as an “attractor model,” since in the absence of noise the deterministic dynamics asymptotically approaches one of two stable fixed points, as in Fig. 12(c). This is in contrast with “oscillator models” in which the attractor of the deterministic component of the dynamics consists of a stable limit cycle, as in Fig. 6 (left column). However, our results show that we can move our model in a smooth, continuous fashion between the two types of dynamics by changing a single parameter. Approaching the bifurcation from one side we can make the period of deterministic oscillation arbitrarily long, so the system spends more and more time in each state. Approaching the bifurcation from the other side, the eigenvalue closest to zero of the stable fixed point can be made arbitrarily close to zero. These smooth changes in the macroscopic dynamics were also noted by Shpiro et al. (2009) who recently concluded from a study of stochastic mean-field-like models that, in order to reproduce observed results, their models must operate with a balance of noise strength

and the strength of the slow processes responsible for switching.

5 Adding noise to the fine-scale model

As mentioned above, there is debate about the role of noise in binocular rivalry (Moreno-Bote et al. 2007; Shpiro et al. 2009). To investigate the role of noise in the fine-scale model we set $B = 1.3$ and added noise to the system by injecting current pulses of the form $I(t) = \pm\beta H(t)e^{-t/8}$ into each excitatory neuron, where H is the Heaviside step function, β is a parameter and t is measured in milliseconds. The arrival times of these pulses form a Poisson process with rate 20 Hz, there is no correlation between arrival times for different neurons, and \pm was chosen with equal probability. Typical resulting dynamics of Φ and χ for $\beta = 2/3$ are shown in Fig. 13, and in Figs. 14 and 15 we show how our estimates of \mathbf{f} and G , respectively, change as the





Fig. 13 The dynamics of Φ and χ with noise of amplitude $\beta = 2/3$ added to the fine-scale model, as described in Section 5. **(a)**: Φ and χ as functions of time. **(b)**: Motion in the Φ, χ plane for a simulation of length 30 s. Motion is in the clockwise direction

Fig. 14 Estimates of f_1 (left column) and f_2 (right column) for noise intensities $\beta = 0$ (top row), $\beta = 2/3$ (second row), $\beta = 4/3$ (third row) and $\beta = 2$ (bottom row)



noise intensity β is increased. We see that there are no qualitative changes in the functions \mathbf{f} and G as β is varied between 0 and 2, but increasing noise level causes the system to explore a larger area of phase space (and spend more time near the origin) and also “smoothes out” the functions \mathbf{f} and G .


In terms of the macroscopic dynamics of the system, the behaviour seen in Fig. 2 and Fig. 13 is qualitatively the same, even though Fig. 2 shows the dynamics of a purely deterministic fine-scale system, while Fig. 13 shows the dynamics of a stochastic fine-scale system. Thus, from the macroscopic point of view, we should regard both systems as having an effective deterministic component and an effective stochastic component. The addition of noise to the fine-scale simulation changes

both the deterministic and the stochastic components of the macroscopic dynamics. Noise intensity can thus be considered as an additional parameter that can vary smoothly through zero.

6 Initialising the system at specific values of the coarse variables

All of the results shown so far have relied on the processing of data from long simulations. Thus we have a great deal of data from frequently-visited areas of phase space but little or none for other regions, leading to poor or non-existent estimates of \mathbf{f} and G —see Fig. 6(b) for example. To overcome this we would like

Fig. 15 Estimates of G_{11} (left column), G_{21} (middle column) and G_{22} (right column) for noise intensities $\beta = 0$ (top row), $\beta = 2/3$ (second row), $\beta = 4/3$ (third row) and $\beta = 2$ (bottom row), in the same format as Fig. 14. The colour ranges are—left column: black= 0, white= 0.05; middle column: black=−0.005, white= 0; right column: black= 0, white= 0.0003



to be able to repeatedly initialise the system at a fine-scale point consistent with a particular macroscopic phase space point and simulate it for a short amount of time, recording the evolution of the macroscopic variables in order to estimate \mathbf{f} and G (Laing et al. 2007). Here we describe a method for doing this and present some results.

Suppose we wish to initialise the system at some “target” point for which $(\chi, \Phi) = (\chi^{\text{targ}}, \Phi^{\text{targ}})$. In what should be considered an unphysical, but computationally useful “preparation step,” we introduce a potential

$$\Lambda(\chi, \Phi) = \alpha_1 (\chi - \chi^{\text{targ}})^2 + \alpha_2 (\Phi - \Phi^{\text{targ}})^2 \quad (29)$$

where $\alpha_1 = 10$ and $\alpha_2 = 350$, and effectively add minus the gradient of this potential to the dynamics for χ and

Φ in order to drive these variables to their target values. Recalling that

$$\chi = \sum_{i=31}^{60} [Ca]^i - \sum_{i=1}^{30} [Ca]^i \quad \text{and} \quad \Phi = \sum_{i=31}^{60} \phi^i - \sum_{i=1}^{30} \phi^i \quad (30)$$

where $[Ca]^i$ is the value of $[Ca]$ for the i th excitatory neuron and similarly for ϕ^i , we add $-\partial\Lambda/\partial[Ca]^i$ to the right hand side of the existing equation for $d[Ca]^i/dt$ and add $-\partial\Lambda/\partial\phi^i$ to the right hand side of the existing equation for $d\phi^i/dt$ (see Eqs. (35) and (36) in Appendix), to create what we refer to as the “constrained” system. To approximately initialise the system at $(\chi, \Phi) = (\chi^{\text{targ}}, \Phi^{\text{targ}})$ we run the constrained system for a short amount of time, using as an initial condition a data-point from an *unconstrained* simulation with values of (χ, Φ) close to $(\chi^{\text{targ}}, \Phi^{\text{targ}})$. If, after this short constrained simulation, the values of




Fig. 16 Blue (black) arrows: deterministic vector field as estimated using the techniques in Section 3.1 (Section 6). Red (green) curve: trajectory following the deterministic vector field shown with blue (black) arrows

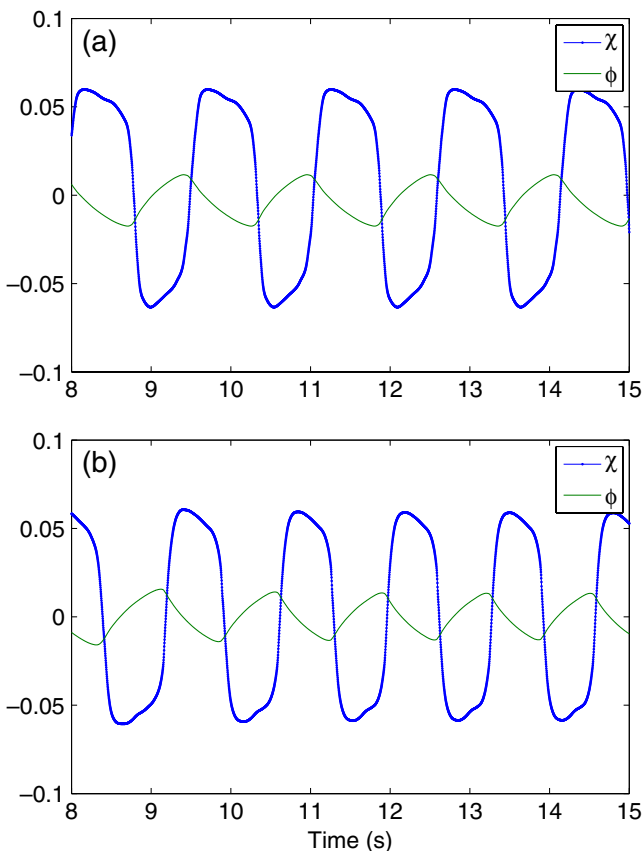


Fig. 17 (a): χ and Φ as functions of time for the trajectory shown in green in Fig. 16, i.e. using the vector field estimated by initialising based on coarse (χ, Φ) values. (b): χ and Φ as functions of time for the trajectory shown in red in Fig. 16, i.e. using the vector field estimated using a long simulation of the fine-scale model

(χ, Φ) are close enough to $(\chi^{\text{targ}}, \Phi^{\text{targ}})$, the constraint is removed, the “natural” system is simulated for a further 10 ms and \mathbf{f} and G are estimated as in Section 3.1. If the values of (χ, Φ) are not close enough to $(\chi^{\text{targ}}, \Phi^{\text{targ}})$, this initial condition is rejected and another is chosen. This procedure is repeated a number of times for a specific $(\chi^{\text{targ}}, \Phi^{\text{targ}})$, and the values of $(\chi^{\text{targ}}, \Phi^{\text{targ}})$ are taken from a uniform rectangular grid.

As a demonstration of our results, Fig. 16 shows, with blue arrows, the deterministic vector field as estimated in Section 3.1, overlaid (in red) with a simulation of this field, as in Fig. 6(b). Also shown, with black arrows, is the deterministic vector field as estimated using the initialisation procedure discussed immediately above, on a 15×15 grid and (in green) a simulation using this vector field. The two attracting limit cycles are quite close and, given the sparseness of the grid upon which we have estimates of the vector field, the agreement is quite good. Figure 17 (top/bottom) shows the time series of χ and Φ associated with the green/red curve in Fig. 16.

7 Conclusion

We have presented several principled derivations of macroscopic models for binocular rivalry from a previously-studied fine-scale model (Laing and Chow 2002). We used both our experience and data-mining tools to extract several appropriate variables, and then processed the results of simulations to determine functions governing the dynamics of these variables. The behaviour of the reduced models agreed very well with that of the fine-scale model. We also discussed parameter variation, the role of noise, and a method for initialising the coarse variables.

We now discuss more generally the types of results that can be obtained using the techniques demonstrated here and their significance. As mentioned in Section 3.1, for fixed parameters it is much quicker to simulate a Langevin equation using, for example, the scheme in Eq. (13) than to simulate the fine-scale model. However, much more can be obtained from the estimates of \mathbf{f} and G . For example, by using interpolation in phase space on the components of \mathbf{f} , the position of the unstable fixed point that appears to be inside the limit cycle of Fig. 6(b) could be found. Normal simulations of the fine-scale model cannot be used to find this fixed point, since it is unstable. As shown in Section 6, the fact that the unconstrained fine-scale model spends little time near this unstable fixed point can be overcome. Alternatively, one may

want to “control” the system, i.e. be able to apply a signal in order to, for example, switch the system from one percept to the other. The reduced models we have derived could be used to design such a controller, which would act based on values of the variables in these models.

As seen in Section 4, bifurcations can occur when parameters are changed. These bifurcations often create or destroy “coarse” unstable invariant sets (fixed points, periodic orbits). These unstable sets will not be observed simply by simulating the fine-scale model, but can be found using the type of reduced model presented here. It is this fact that enables us to say that the bifurcation seen in Section 4 is a symmetric SNIC bifurcation rather than, say, a bifurcation involving a heteroclinic connection between two symmetry-related saddle fixed points. We can also use this approach to determine the precise value of a parameter at which the system switches from being an “oscillator model” to being an “attractor model” (Moreno-Bote et al. 2007).

Other interesting observations include the fact that the noise in the presumed Langevin equation is multiplicative, i.e. the components of the matrix G vary as \mathbf{X} is varied—see Fig. 5. Also, the fact that only two stochastic differential equations (for χ and Φ , or ν_1 and ν_2) are required to reproduce the dynamics of the fine-scale system is consistent with the smallest number of variables that an “oscillator model” (Moreno-Bote et al. 2007) must have. (At least for the parameter values we considered. The number of variables required in a reduced model may well change as parameters are varied. Ideally the number of variables used should change adaptively—by analogy with variable step-size in adaptive numerical integrators—as parameters are varied, in a way that preserves the accuracy of the reduced model (Makeev et al. 2002)).

The problem of deriving accurate “coarse” scale models from “fine” scale models is one of the outstanding problems in computational neuroscience (see Ermentrout 1994; Shriki et al. 2003, for several successes). Although the results shown here are computationally intensive, the method does provide an alternative to the special cases just mentioned, for which analytical progress is possible. The techniques are very general, and should be applicable to a number of other systems which show “emergent” macroscopic behaviour, and for which we have detailed biophysical models (Marder and Bucher 2001; Rybak et al. 2004).

Acknowledgements The work of C.R.L. was partially supported by the Marsden Fund, administered by The Royal Society of New Zealand. The work of I.G.K. and T.F. was partially supported by the National Science Foundation and DARPA.

Appendix: A model equations

Here we present the model equations. They are very similar to those in Laing and Chow (2002). For each excitatory neuron we have

$$\frac{dV_e}{dt} = I_{syn} + I_{ext} - I_{mem}(V_e, n_e, h_e) - I_{AHP} \quad (31)$$

$$\frac{dn_e}{dt} = \psi[\alpha_n(V_e)(1 - n_e) - \beta_n(V_e)n_e] \quad (32)$$

$$\frac{dh_e}{dt} = \psi[\alpha_h(V_e)(1 - h_e) - \beta_h(V_e)h_e] \quad (33)$$

$$\tau_e \frac{ds_e}{dt} = A\sigma(V_e)(1 - s_e) - s_e \quad (34)$$

$$\frac{d[Ca]}{dt} = \frac{-0.002g_{Ca}(V_e - V_{Ca})}{1 + \exp[-(V_e + 25)/2.5]} - [Ca]/80 \quad (35)$$

$$\tau_g \frac{d\phi}{dt} = 1 - \phi - B\sigma(V_e)\phi \quad (36)$$

where $I_{mem}(V_e, n_e, h_e) = g_L(V_e - V_L) + g_K n_e^4(V_e - V_K) + g_{Na}(m_\infty(V_e))^3 h_e(V_e - V_{Na})$ and $I_{AHP} = g_{AHP}[Ca]/([Ca] + 1)(V_e - V_K)$. Other functions are $m_\infty(V) = \alpha_m(V)/(\alpha_m(V) + \beta_m(V))$, $\alpha_m(V) = 0.1(V + 30)/(1 - \exp[-0.1(V + 30)])$, $\beta_m(V) = 4 \exp[-(V + 55)/18]$, $\alpha_n(V) = 0.01(V + 34)/(1 - \exp[-0.1(V + 34)])$, $\beta_n(V) = 0.125 \exp[-(V + 44)/80]$, $\alpha_h(V) = 0.07 \exp[-(V + 44)/20]$, $\beta_h(V) = 1/(1 + \exp[-0.1(V + 14)])$, $\sigma(V) = 1/(1 + \exp[-(V + 20)/4])$. Parameters are $g_L = 0.05$ mS/cm², $V_L = -65$ mV, $g_K = 40$ mS/cm², $V_K = -80$ mV, $g_{Na} = 100$ mS/cm², $V_{Na} = 55$ mV, $V_{Ca} = 120$ mV, $g_{AHP} = 0.05$ mS/cm², $g_{Ca} = 0.1$ mS/cm², $\psi = 3$, $\tau_e = 8$ ms, $\tau_g = 1,000$ ms and $A = 20$. B is initially 1.3, but is varied in Section 4.

For each inhibitory neuron we have

$$\frac{dV_i}{dt} = I_{syn} - I_{mem}(V_i, n_i, h_i) \quad (37)$$

$$\frac{dn_i}{dt} = \psi[\alpha_n(V_i)(1 - n_i) - \beta_n(V_i)n_i] \quad (38)$$

$$\frac{dh_i}{dt} = \psi[\alpha_h(V_i)(1 - h_i) - \beta_h(V_i)h_i] \quad (39)$$

$$\tau_i \frac{ds_i}{dt} = A\sigma(V_i)(1 - s_i) - s_i \quad (40)$$

where $\tau_i = 10$ ms and other functions are as above.

The synaptic current entering the j th excitatory neuron is

$$(V_+ - V_e^j) \frac{1}{N} \sum_{k=1}^N g_{ee}^{jk} s_e^k \phi^k + (V_- - V_e^j) \frac{1}{N} \sum_{k=1}^N g_{ie}^{jk} s_i^k \quad (41)$$

where V_e^j is the voltage of the j th excitatory neuron in mV, $s_{e/i}^k$ is the strength of the synapse emanating

from the k th excitatory/inhibitory neuron, ϕ^k is the factor by which the k th excitatory neuron is depressed, $N = 60$ is the number of excitatory neurons (equal to the number of inhibitory neurons), and the Gaussian coupling functions are given by

$$g_{ee}^{jk} = \alpha_{ee} \sqrt{50/\pi} \exp \{-50[(j - k)/N]^2\} \quad (42)$$

and

$$g_{ie}^{jk} = \alpha_{ie} \sqrt{20/\pi} \exp \{-20[(j - k)/N]^2\}. \quad (43)$$

The reversal potentials are $V_+ = 0$ mV, $V_- = -80$ mV. Similarly, the synaptic current entering the j th inhibitory neuron is

$$(V_+ - V_i^j) \frac{1}{N} \sum_{k=1}^N g_{ei}^{jk} s_e^k + (V_- - V_i^j) \frac{1}{N} \sum_{k=1}^N g_{ii}^{jk} s_i^k \quad (44)$$

where V_i^j is the voltage of the j th inhibitory neuron in mV, and the coupling functions are

$$g_{ei}^{jk} = \alpha_{ei} \sqrt{20/\pi} \exp \{-20[(j - k)/N]^2\} \quad (45)$$

and

$$g_{ii}^{jk} = \alpha_{ii} \sqrt{30/\pi} \exp \{-30[(j - k)/N]^2\}. \quad (46)$$

Parameters are $\alpha_{ee} = 0.285$ mS/cm², $\alpha_{ie} = 0.36$ mS/cm², $\alpha_{ei} = 0.2$ mS/cm² and $\alpha_{ii} = 0.07$ mS/cm². The external current to the j th excitatory neuron in μ A/cm² is

$$I_{ext}^j = \frac{0.4}{\sqrt{2}} \left[\exp \left(- \left[\frac{10[j - N/4]}{N} \right]^2 \right) + \exp \left(- \left[\frac{10[j - 3N/4]}{N} \right]^2 \right) \right] - 0.01 \quad (47)$$

i.e. the external current injected into the excitatory population consists of two Gaussians, centered at 1/4 and 3/4 of the way around the domain. The equations were simulated using Euler’s method with a fixed time-step of 0.02 ms, and no significant changes in the network behaviour were observed when time-steps of 0.01 or 0.005 ms were used.

Note that the system is completely deterministic. Section 5 discusses the results from simulating a stochastic version of this network.

References

Ashwin, P., & Lavric, A. (2010). A low-dimensional model of binocular rivalry using winnerless competition. *Physica D* (in press).
 Bahraminasab, A., Kenwright, D., Stefanovska, A., Ghasemi, F., & McClintock, P. (2008). Phase coupling in the cardiorespiratory interaction. *IET Systems Biology*, 2, 48–54.

Bengio, Y., Delalleau, O., Le Roux, N., Paiement, J., Vincent, P., & Ouimet, M. (2004). Learning eigenfunctions links spectral embedding and kernel PCA. *Neural Computation*, 16(10), 2197–2219.
 Blake, R. (2001). A primer on binocular rivalry, including current controversies. *Brain and Mind*, 2(1), 5–38.
 Blake, R., & Logothetis, N. (2002). Visual competition. *Nature Reviews Neuroscience*, 3(1), 1–11.
 Cai, D., Tao, L., Shelley, M., & McLaughlin, D. (2004). An effective kinetic representation of fluctuation-driven neuronal networks with application to simple and complex cells in visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 101(20), 7757.
 Coifman, R., & Lafon, S. (2006). Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1), 5–30.
 Dayan, P. (1998). A hierarchical model of binocular rivalry. *Neural Computation*, 10(5), 1119–1135.
 Erban, R., Frewen, T., Wang, X., Elston, T., Coifman, R., Nadler, B., et al. (2007). Variable-free exploration of stochastic models: A gene regulatory network example. *The Journal of Chemical Physics*, 126, 155,103.
 Ermentrout, B. (1994). Reduction of conductance-based models with slow synapses to neural nets. *Neural Computation*, 6(4), 679–695.
 Freeman, A. (2005). Multistage model for binocular rivalry. *Journal of Neurophysiology*, 94(6), 4412–4420.
 Friedrich, R., Siegert, S., Peinke, J., Lück, S., Siefert, M., Lindemann, M., et al. (2000). Extracting model equations from experimental data. *Physics Letters A*, 271(3), 217–222.
 Gerstner, W., & Kistler, W. (2002). *Spiking neuron models: An introduction*. New York: Cambridge University Press.
 Gradišek, J., Siegert, S., Friedrich, R., & Grabec, I. (2000). Analysis of time series from stochastic processes. *Physical Review E*, 62(3), 3146–3155.
 Grossberg, S., Yazdanbakhsh, A., Cao, Y., & Swaminathan, G. (2008). How does binocular rivalry emerge from cortical mechanisms of 3-D vision? *Vision Research*, 48(21), 2232–2250.
 Gutkin, B., Laing, C., Colby, C., Chow, C., & Ermentrout, G. (2001). Turning on and off with excitation: The role of spike-timing asynchrony and synchrony in sustained neural activity. *Journal of Computational Neuroscience*, 11(2), 121–134.
 Jolliffe, I. (2002). *Principal component analysis*. New York: Springer.
 Jolly, M., Kevrekidis, I., & Titi, E. (1990). Approximate inertial manifolds for the Kuramoto–Sivashinsky equation: Analysis and computations. *Physica D*, 44(1–2), 38–60.
 Kalarickal, G., & Marshall, J. (2000). Neural model of temporal and stochastic properties of binocular rivalry. *Neurocomputing*, 32–33, 843–853.
 Kevrekidis, I., Gear, C., Hyman, J., Kevrekidis, P., Runborg, O., & Theodoropoulos, C. (2003). Equation-free, coarse-grained multiscale computation: Enabling microscopic simulators to perform system-level analysis. *Communications in Mathematical Sciences*, 1(4), 715–762.
 Kuusela, T., Shepherd, T., & Hietarinta, J. (2003). Stochastic model for heart-rate fluctuations. *Physical Review E*, 67(6), 061,904.
 Lago-Fernandez, L., & Deco, G. (2002). A model of binocular rivalry based on competition in IT. *Neurocomputing*, 44–46, 503–507.
 Laing, C. (2006). On the application of “equation-free” modelling to neural systems. *Journal of Computational Neuroscience*, 20(1), 5–23.

- Laing, C., & Chow, C. (2001). Stationary bumps in networks of spiking neurons. *Neural Computation*, *13*(7), 1473–1494.
- Laing, C., & Chow, C. (2002). A spiking neuron model for binocular rivalry. *Journal of Computational Neuroscience*, *12*(1), 39–53.
- Laing, C., Frewen, T., & Kevrekidis, I. (2007). Coarse-grained dynamics of an activity bump in a neural field model. *Nonlinearity*, *20*(9), 2127–2146.
- Leopold, D., & Logothetis, N. (1999). Multistable phenomena: Changing views in perception. *Trends in Cognitive Sciences*, *3*(7), 254–264.
- Logothetis, N., Leopold, D., & Sheinberg, D. (1996). What is rivaling during binocular rivalry? *Nature*, *380*(6575), 621–624.
- Makeev, A., Maroudas, D., & Kevrekidis, I. (2002). “Coarse” stability and bifurcation analysis using stochastic simulators: Kinetic Monte Carlo examples. *The Journal of Chemical Physics*, *116*, 10,083.
- Marder, E., & Bucher, D. (2001). Central pattern generators and the control of rhythmic movements. *Current Biology*, *11*(23), 986–996.
- Moreno-Bote, R., Rinzel, J., & Rubin, N. (2007). Noise-induced alternations in an attractor network model of perceptual bistability. *Journal of Neurophysiology*, *98*(3), 1125.
- van Mourik, A., Daffertshofer, A., & Beek, P. (2006). Deterministic and stochastic features of rhythmic human movement. *Biological Cybernetics*, *94*(3), 233–244.
- Nadler, B., Lafon, S., Coifman, R., & Kevrekidis, I. (2006). Diffusion maps, spectral clustering and reaction coordinates of dynamical systems. *Applied and Computational Harmonic Analysis*, *21*, 113–127.
- Ragwitz, M., & Kantz, H. (2001). Indispensable finite time corrections for Fokker–Planck equations from time series data. *Physical Review Letters*, *87*(25), 254,501.
- Rega, G., & Troger, H. (2005). Dimension reduction of dynamical systems: Methods, models, applications. *Nonlinear Dynamics*, *41*(1), 1–15.
- Rybak, I., Shevtsova, N., Paton, J., Dick, T., St-John, W., Mörschel, M., et al. (2004). Modeling the ponto-medullary respiratory network. *Respiratory Physiology & Neurobiology*, *143*(2–3), 307–319.
- Shpiro, A., Moreno-Bote, R., Rubin, N., & Rinzel, J. (2009). Balance between noise and adaptation in competition models of perceptual bistability. *Journal of Computational Neuroscience*, *27*(1), 37–54.
- Shriki, O., Hansel, D., & Sompolinsky, H. (2003). Rate models for conductance-based cortical neuronal networks. *Neural Computation*, *15*(8), 1809–1841.
- Stollenwerk, L., & Bode, M. (2003). Lateral neural model of binocular rivalry. *Neural Computation*, *15*(12), 2863–2882.
- Tong, F., Meng, M., & Blake, R. (2006). Neural bases of binocular rivalry. *Trends in Cognitive Sciences*, *10*(11), 502–511.
- Tranchina, D. (2009). Population density methods in large-scale neural network modelling. In C. Laing, & G. J. Lord (eds) *Stochastic methods in neuroscience* (pp. 181–216). Oxford: Oxford University Press.
- Wilson, H. (2003). Computational evidence for a rivalry hierarchy in vision. *Proceedings of the National Academy of Sciences, USA*, *100*, 14,499–14,503.
- Wilson, H., & Cowan, J. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*(1), 1–24.