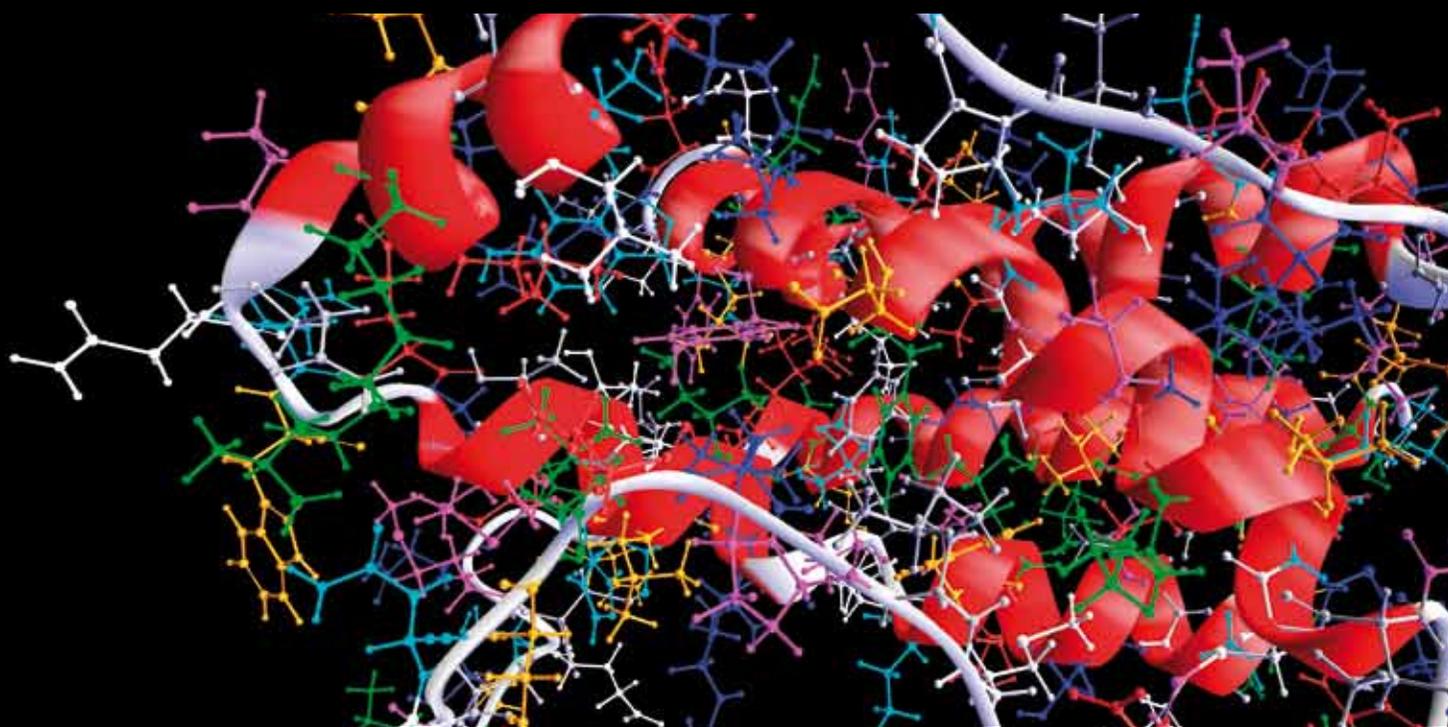


NONLINEAR PHENOMENA IN BIOLOGY AND MEDICINE

GUEST EDITORS: VIKAS RAI, SREENIVASAN R. NADAR, T. R. KRISHNAMOHAN,
AND RANJIT KUMAR UPADHYAY





Nonlinear Phenomena in Biology and Medicine

Computational and Mathematical Methods in Medicine

Nonlinear Phenomena in Biology and Medicine

Guest Editors: Vikas Rai, Sreenivasan R. Nadar,
T. R. Krishnamohan, and Ranjit Kumar Upadhyay



Copyright © 2012 Hindawi Publishing Corporation. All rights reserved.

This is a special issue published in “Computational and Mathematical Methods in Medicine.” All articles are open access articles distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Editorial Board

Zvia Agur, Israel
Emil Alexov, USA
Gary C. An, USA
Georgios Archontis, Cyprus
Pascal Auffinger, France
Facundo Ballester, Spain
Dimos Baltas, Germany
Izhar Bar-Gad, Israel
Chris Bauch, Canada
Maxim Bazhenov, USA
Niko Beerenwinkel, Switzerland
Philip Biggin, UK
Michael Breakspear, Australia
Thierry Busso, France
Bill Crum, UK
Timothy David, New Zealand
Gustavo Deco, Spain
Carmen Domene, UK
Frank Emmert-Streib, UK
Alfonso T. García-Sosa, Estonia
Kannan Gunasekaran, USA
Damien R. Hall, Japan

William F. Harris, South Africa
Vassily Hatzimanikatis, USA
Volkhard Helms, Germany
J.-H. S. Hofmeyr, South Africa
Seiya Imoto, Japan
Bleddyn Jones, UK
Lawrence A. Kelley, UK
Lev Klebanov, Czech Republic
Ina Koch, Germany
David Liley, Australia
Quan Long, UK
Yoram Louzoun, Israel
C.-M. C. Ma, USA
Jianpeng Ma, USA
Reinoud Maex, France
Francois Major, Canada
Ramit Mehr, Israel
Seth Michelson, USA
Michele Migliore, Italy
Karol Miller, Australia
Ernst Niebur, USA
Kazuhisa Nishizawa, Japan

Martin Nowak, USA
Markus Owen, UK
Hugo Palmans, UK
Lech S. Papiez, USA
Jean Pierre Rospars, France
David James Sherman, France
Sivabal Sivaloganathan, Canada
Stephen W. Smye, UK
Merryn H. Tawhai, New Zealand
Elisabeth Tillier, Canada
Nestor V. Torres, Spain
Anna Tramontano, Italy
Nelson J. Trujillo-Barreto, Cuba
Kutlu O. Ulgen, Turkey
Nagarajan Vaidehi, USA
Edelmira Valero, Spain
Wim van Drongelen, USA
Jinliang Wang, UK
Jacek Waniewski, Poland
Guang Wu, China
Henggui Zhang, UK

Contents

Nonlinear Phenomena in Biology and Medicine, Vikas Rai, Sreenivasan R. Nadar, T. R. Krishnamohan, and Ranjit Kumar Upadhyay
Volume 2012, Article ID 159589, 2 pages

Inference for Ecological Dynamical Systems: A Case Study of Two Endemic Diseases, Daniel A. Vasco
Volume 2012, Article ID 390694, 13 pages

Higuchi Fractal Properties of Onset Epilepsy Electroencephalogram, Truong Quang Dang Khoa, Vo Quang Ha, and Vo Van Toi
Volume 2012, Article ID 461426, 6 pages

Linear and Nonlinear Heart Rate Variability Indexes in Clinical Practice, Buccelletti Francesco, Bocci Maria Grazia, Gilardi Emanuele, Fiore Valentina, Calcinaro Sara, Fragnoli Chiara, Maviglia Riccardo, and Franceschi Francesco
Volume 2012, Article ID 219080, 5 pages

A Finite Element Study of Micropipette Aspiration of Single Cells: Effect of Compressibility, Amirhossein Jafari Bidhendi and Rami K. Korhonen
Volume 2012, Article ID 192618, 9 pages

Higher-Order Spectrum in Understanding Nonlinearity in EEG Rhythms, Cauchy Pradhan, Susant K. Jena, Sreenivasan R. Nadar, and N. Pradhan
Volume 2012, Article ID 206857, 8 pages

Simulation of Spread and Control of Lesions in Brain, Krishna Mohan Thamattoor Raman
Volume 2012, Article ID 383546, 6 pages

Free Energy, Value, and Attractors, Karl Friston and Ping Ao
Volume 2012, Article ID 937860, 27 pages

Editorial

Nonlinear Phenomena in Biology and Medicine

Vikas Rai,¹ Sreenivasan R. Nadar,² and Ranjit K. Upadhyay³

¹Department of Mathematics, Faculty of Science, Jazan University, 45142 Jizan, Saudi Arabia

²National Institute of Mental Health, Bethesda, 20892 MD, USA

³Indian School of Mines, 826004 Dhanbad, India

Correspondence should be addressed to Vikas Rai, rvikas41@hotmail.com

Received 12 February 2012; Accepted 12 February 2012

Copyright © 2012 Vikas Rai et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Nonlinear dynamics has changed our view of many biological phenomena; e.g., coexistence of consumer species competing for the same resource [1], biological pattern formation; e.g., clustering and differentiation of mesenchymal cells during pre-cartilage mesenchymal condensation in chick limb chondrogenesis (bone patterning and formation) [2] and analysis of heart rate variability in healthy as well as diseased subjects [3] to name a few. The diversity of disciplines and successes of approaches based on nonlinear dynamics, complexity theory, and systems biology in resolving a few difficult issues and answering some outstanding questions in recent years led to the idea of compiling a special volume on the subject.

This special issue is dedicated to nonlinear phenomena in biology and medicine. The research papers which appear in this issue can be classified into two categories: (1) papers which discuss approaches to nonlinear time series analysis based on dynamical systems theory and (2) others which present models and their potential application to human behavior, biology and medicine.

In this issue, K. Friston and F. Ao draw attention of researchers in neuroscience to the possibility that action and perception can be understood in terms of minimization of a free energy surface functional which minimizes “sensory surprise”. Free energy is an information theoretic measure that bounds the surprise on sampling data, given a generative model. Authors compare two different approaches to describe agent action and prediction based on the free energy principle and optimal control reinforcement learning. In the first case, an agent’s behavior is controlled by a free energy surface functional that minimizes “sensory surprise” along the trajectory in the phase space of a nonautonomous system with strong fluctuations. The mathematical definition

of surprise is conditional entropy. In the last case, an agent’s adaptive behavior is determined in order to maximize a reward.

Epilepsy is the principal brain dysfunction which affects about 1% world population and has important public health implication. The conventional signal analysis methods such as the count of focal spike density, the frequency coherence, or spectral analyses are not reliable predictors. In this issue the paper by T. Khoa et al. describes a method based on fractal dimension as a principal tool to diagnose epilepsy. Their method combines independent component algorithm with averaging filter at the preprocessing step. The authors show that this improved method could be used to analyze EEG signals to diagnose epilepsy.

B. Francesco et al. review the applications of linear and nonlinear indexes in heart rate variability and show how these indexes are useful in clinical practice using data from patients.

D. Vasco presents a contribution which focuses on a Bayesian Markov Chain Monte Carlo method to infer parameters in simple SIR model in epidemiology. Application of Markov chain models in conjunction with SSA is not new. The interesting part of the paper is that the author has successfully employed this method to measles and pertussis epidemic time series data from 60 UK cities.

K. Mohan presents a simulation model to understand the spread and control of lesions based on a planar graph representation for the central nervous system. The author demonstrates that the model is capable of generating a wide variety of lesion growth and arrest scenarios.

A. Bidhendi and R. Korhonen contribute a paper on nonlinear deformations of cells during micropipette aspiration procedure to measure its viscoelastic properties. The

paper examines a neo-Hookean model for micropipette aspiration. A fillet radius is considered at the opening of the micropipette to study its effect on the modeled response of the cell. The authors estimate optimal parameters of the model from the experimental stem cell data. Their findings suggest that the compressibility and bulk relaxation/fluid flow play a significant role in the deformation behavior of single cells and should be taken into account in the analysis of the mechanics of cells.

C. Pradhan et al. examine the fundamental nature of the brain electrical activities recorded as electroencephalogram (EEG). Linear stochastic models and spectral estimates are the most common methods for the analysis of EEG because of their robustness, simplicity of interpretation, and apparent association with rhythmic behavioral patterns in nature. The paper extends the application of higher-order spectrum in order to clarify the hidden characteristics of EEG signals that simply do not arise of random processes. This paper demonstrates the suitability of bispectral analysis to distinguish chaotic systems from filtered noises and normal background EEG activity.

In sum, we note that these contributions present state of the art of their respective subdisciplines. The research papers appearing in this special issue will serve as a guide to what is yet to follow in this fascinating field of biology and medicine. Procedures laid down in the research papers by T. Khoa et al. and by C. Pradhan et al. can be combined to design a protocol for the diagnostics of epileptic disorders. We hope that this volume will serve interests of researchers working in the field of applied biology and medicine.

Vikas Rai
Sreenivasan R. Nadar
Ranjit K. Upadhyay

References

- [1] P. A. Abrams, C. E. Brassil, and R. D. Holt, "Dynamics and responses to mortality rates of competing predators undergoing predator-prey cycles," *Theoretical Population Biology*, vol. 64, no. 2, pp. 163–176, 2003.
- [2] W. Zeng, G. L. Thomas, and J. A. Glazier, "Non-turing stripes and spots: a mechanism for biological cell clustering," *Physica A*, vol. 341, pp. 482–494, 2004.
- [3] A. L. Goldberger, L. A. N. Amaral, J. M. Hausdorff, P. Ch. Ivanov, C.-K. Peng, and H. E. Stanley, "Fractal dynamics in physiology: alterations with disease and aging," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, no. 1, supplement, pp. 2466–2472, 2002.

Research Article

Inference for Ecological Dynamical Systems: A Case Study of Two Endemic Diseases

Daniel A. Vasco

Department of Biology, Duke University, Box 90338, Durham, NC 27708, USA

Correspondence should be addressed to Daniel A. Vasco, vasco.daniel.a@gmail.com

Received 2 September 2011; Revised 19 November 2011; Accepted 21 November 2011

Academic Editor: Vikas Rai

Copyright © 2012 Daniel A. Vasco. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

A Bayesian Markov chain Monte Carlo method is used to infer parameters for an open stochastic epidemiological model: the Markovian susceptible-infected-recovered (SIR) model, which is suitable for modeling and simulating recurrent epidemics. This allows exploring two major problems of inference appearing in many mechanistic population models. First, trajectories of these processes are often only partly observed. For example, during an epidemic the transmission process is only partly observable: one cannot record infection times. Therefore, one only records cases (infections) as the observations. As a result some means of imputing or reconstructing individuals in the susceptible cases class must be accomplished. Second, the official reporting of observations (cases in epidemiology) is typically done not as they are actually recorded but at some temporal interval over which they have been aggregated. To address these issues, this paper investigates the following problems. Parameter inference for a perfectly sampled open Markovian SIR is first considered. Next inference for an imperfectly observed sample path of the system is studied. Although this second problem has been solved for the case of closed epidemics, it has proven quite difficult for the case of open recurrent epidemics. Lastly, application of the statistical theory is made to measles and pertussis epidemic time series data from 60 UK cities.

1. Introduction

The linking of ecological theory with data is currently a major scientific challenge. Modern methods of data collection and storage are rapidly improving at all levels, from the detailed study of individuals in populations to the distribution of populations and communities over vast landscapes. Despite the ease with which it is possible to develop statistical theory and Bayesian Markov chain Monte Carlo (MCMC) computational statistics for many ecological problems [1], the resolution of many computational issues for these problems remains largely unresolved when fitting dynamical ecological models (either in discrete or continuous time) to large ecological and public health data sets.

In fact, it is possible to discuss many of these computational difficulties using simple stochastic epidemiological models. Epidemiological processes serve as excellent prototypes for exhibiting two major problems of inference that appear in many mechanistic dynamic models. First, the transmission process during an epidemic is only partly

observed. As a result in epidemiology one only records cases and rarely observes the infection time precisely. Second, the official reporting of observations (cases in epidemiology) is typically done not as they are actually recorded, but at some temporal interval over which they have been aggregated. Although these problems have largely been solved for the case of closed epidemics, it has proven quite difficult for the case of open populations that produce recurrent epidemics (endemic diseases) over many generations in continuous time. This is because it is hard to simulate paths that are consistent with the data due to the condition that one must sample from many recorded intervals given the number of infectives in the beginning and at the end of the interval. In general this has proven easier to do for short-duration epidemics because of computational limitations due to data augmentation. As the number of recorded intervals increases the data likelihood computation rapidly becomes intractable or impossible.

In this paper a data augmentation strategy is implemented that allows addressing these problems, is reasonably

straightforward to implement, is fast and accurate for the problem at hand. The basis of the method is a recently proposed Bayesian MCMC algorithm proposed by Wilkinson [2]. This algorithm is used as the computational foundation for inferring parameters using a stochastic epidemiological model: the Markovian susceptible-infected-recovered (SIR) model of epidemiology. The approach used here includes births and deaths as well as immigration of infectives and hence allows modeling of recurrent epidemics and the inference of model parameters for endemic diseases. The computational methods in this paper are largely drawn from recent approaches taken in systems biology for inference of parameters using time series data. In the Results and Discussion (Section 3.4) part of this paper a brief review is made of these computational methodologies. They are compared to the Bayesian approach taken here along with its advantages and limitations.

Most previous work on the SIR using likelihood [3] and Bayesian MCMC [4, 5] has focused on epidemic data sets collected in small closed communities such as households [6, 7] but very little into endemic diseases [8]. Exceptions to this trend are work by Gibson and Renshaw [9] and the more recent work of Cauchemez and Ferguson [8]. The form of the likelihood in the current framework is the same as that presented in O’Neill and Roberts [4] and similar to that of Cauchemez and Ferguson [8] except that in the present study births, deaths, and immigration of infected cases are included in the dynamics. This makes the SIR likelihood used here most similar to that first utilized by Gibson and Renshaw [9]. This assumption is critical in simulating an open population stochastic SIR as an approximate model of endemic diseases. Adding an influx of migrants allows computationally generating patterns of persistent and complex sustained epidemic oscillations [10, 11].

Application of the inference method is made for time series data for two endemic childhood diseases, pertussis and measles. It is shown how to reconstruct stochastic oscillations using simulations and model checking with respect to observed cases. Finally the hypothesis of coherence resonance is investigated and it is shown how it may account for some of the empirically observed patterns of stochastic oscillatory dynamics of the two endemic diseases.

2. Materials and Methods

2.1. SIR Inference: Perfect Information. In this paper a stochastic version of the Kermack-McKendrick susceptible-infectious-recovered (SIR) model [12] will be used to address the inference problem of mechanistic modeling in ecology. As shown below, a structured representation of this model (in fact, any mechanistic model in ecology) can immediately be used to derive a corresponding Markovian stochastic population model. In the deterministic SIR model, there are seven possible events: birth, death (including all possible labeled events for each type of death event), transmission,

recovery, and immigration. The deterministic framework is described by a set of coupled ordinary differential equations:

$$\frac{ds}{dt} = -\frac{\beta}{N}i(t)s(t) + \mu_1N - \mu_2s(t), \quad (1)$$

$$\frac{di}{dt} = -\gamma i(t) + \frac{\beta}{N}i(t)s(t) + \sigma - \mu_3i(t), \quad (2)$$

$$\frac{dr}{dt} = \gamma i(t) - \mu_4r(t) - \sigma. \quad (3)$$

Here, β represents the transmission rate, σ denotes the rate of immigration of infectious individuals, and $1/\gamma$ describes the average infectious period [13]. The immigration term σ is placed in the recovered equation to ensure constant population—a basic assumption underlying the SIR model. The μ_i represent the birth and death rates for each compartment. Note, however, that *per capita* birth and death rates may be assumed to be the same (μ), ensuring a constant long-term population size, N . Also note that here $\mathbf{x} = (x_1(t), x_2(t), x_3(t)) = (s(t), i(t), r(t))$, the 3-dimensional vector of state variables.

Next consider an event-driven model of state change. Define $\alpha = (\beta, \mu_1, \mu_2, \gamma, \mu_3, \sigma, \mu_4)$, which is the 7-dimensional vector of parameters associated with the SI transitions (transitions to the recovered class will be ignored in this paper):

$$\begin{pmatrix} \alpha_1 e_1 \\ \alpha_2 e_2 \\ \alpha_3 e_2 \\ \alpha_4 e_4 \\ \alpha_5 e_5 \\ \alpha_6 e_6 \\ \alpha_7 e_7 \end{pmatrix} = \begin{pmatrix} \text{event 1} \\ \text{event 2} \\ \text{event 3} \\ \text{event 4} \\ \text{event 5} \\ \text{event 6} \\ \text{event 7} \end{pmatrix} = \begin{pmatrix} \text{transmission} \\ \text{birth susceptible} \\ \text{death susceptible} \\ \text{infection} \\ \text{death infected} \\ \text{immigration} \\ \text{death recovered} \end{pmatrix} = \begin{pmatrix} \frac{\beta}{N}si \\ \mu_1N \\ \mu_2s \\ \gamma i \\ \mu_3i \\ \sigma \\ \mu_4r \end{pmatrix}. \quad (4)$$

Define a change in state as if it occurred from some updating rule applied to each possible event. The updating rules are constrained by the structure of the continuous open SIR equations (1)–(4) to specify an association between the event function, $e_j(\mathbf{x})$, and an associated stage-change vector ν_j . Define an *event pathway* vector, $(\mathcal{P}_1, \dots, \mathcal{P}_7)$, where each path \mathcal{P}_j in Table 1 describes a transition event giving rise to integral state changes in S and I . Table 1 shows how events defined by (4) along with the updating of SI states in the SIR model given by (1) and (2) may be used as a template to construct these pathways. The structured representation immediately gives the transition probabilities for the Markovian SIR model [14, 15]. Equations (1)–(3) may now be used to specify the probability event function, $e_j(\mathbf{x})$, and the associated stage-change vector ν_j . For example, since one defines event 1 to represent a transmission event, then $e_1 = P_1(\Delta S^1 = \nu_{11}, \Delta I^1 = \nu_{12})$ has at time t instantaneous rate $e_1(\mathbf{x}, \alpha) = (\beta/N)si$, with $\nu_1 = (\nu_{11}, \nu_{12}) = (-1, 1)$, where P_1 represents the probability of an instantaneous transition for the event path \mathcal{P}_1 . Using the directed network shown in Figure 1, consider event path \mathcal{P}_1 . This is represented in

TABLE 1: Structural representation of SI state-event transitions.

Event path	Parameter	Transition	Flow in node		Flow out node		Flow difference		Change ν_j
	α_j	e_j	S_{in}^j	I_{in}^j	S_{out}^j	I_{out}^j	$\Delta S^j = \nu_{j1}$	$\Delta I^j = \nu_{j2}$	
\mathcal{P}_1	β/N	SI	0	1	1	0	-1	1	ν_1
\mathcal{P}_2	μN	1	1	0	0	0	1	0	ν_2
\mathcal{P}_3	μ	S	0	0	1	0	-1	0	ν_3
\mathcal{P}_4	γ	I	0	0	0	1	0	-1	ν_4
\mathcal{P}_5	μ	I	0	0	0	1	0	-1	ν_5
\mathcal{P}_6	σ	1	0	1	0	0	0	1	ν_6
\mathcal{P}_7	μ	R	0	0	0	0	0	0	ν_7

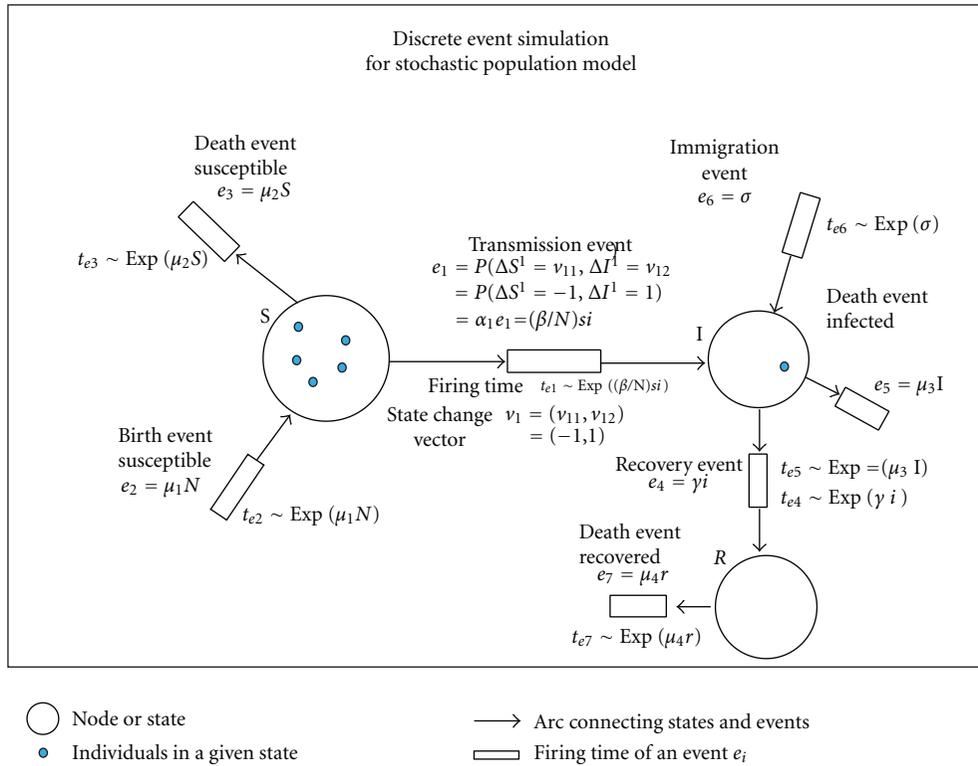


FIGURE 1: Directed network for Markovian SIR dynamics. Each event pathway, \mathcal{P}_j , is represented in the network as the arrow connecting nodes of the random variables, here designated as S , I , and R . The blue dots represent individuals flowing through the network. For example, the flow of individuals out of node with random variable S is represented by the arrow pointing to the box, which gives the firing time to the event and the state-change vector ν_1 associated with the firing of the event.

the network as the arrow connecting nodes of the random S and I variables. The blue dots represent individuals flowing through the network. The flow of individuals out of node with random variable S is represented by the arrow pointing to the box, which gives the firing time to the event and the state-change vector ν_1 associated with the firing of the event. The instantaneous flow (or jumping) between the nodes S and I is determined by $P_1(\Delta S^1 = \nu_{11}, \Delta I^1 = \nu_{12})$. The effect of the firing on an individual in node S is represented by ΔS^1 and corresponds to the first component ν_{11} of the state-change vector ν_1 . The effect on node I is represented by ΔI^1 , which is the second component ν_{12} of state-change vector ν_1 . All of the events in the directed network representing the

open SIR can be treated in the same manner. More generally, one can write

$$e_j(\mathbf{x}, \boldsymbol{\alpha}) dt = \text{probability that } \mathcal{P}_j \text{ event pathway occurs in time } dt. \quad (5)$$

Numerical simulation of stochastic mechanistic models based upon (5) consists of computing the firing of the transitions for each node of the network. The firing of each transition is determined by a random clock running at a time determined by the exponential distribution. For example, in Figure 1 the boxes represent the noisy clocks keeping time

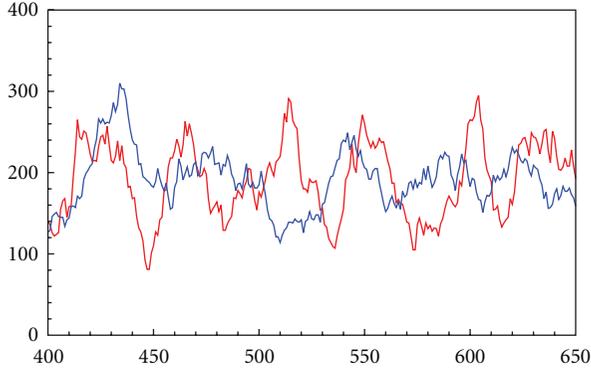


FIGURE 2: A two-hundred-and-fifty-week time series of the number of infected cases from the SIR immigration model discussed in this paper, simulated using the stochastic simulation algorithm (SSA) defined in the supplementary material. y -axis corresponds to observed infected cases. x -axis corresponds to time in weeks. Susceptible cases exhibit a similar complex pattern of noisy oscillations but are not shown. The simulated data are weekly numbers of infected cases; Disease 1 infected cases are shown in blue; Disease 2 infected cases are shown in red. Parameter values used in the simulations: for Disease 1: $\alpha = (\beta, \gamma, \sigma, \mu) = (3.70, .25, .1, .001)$ with $N = 50000$ and for Disease 2: $\alpha = (14.7, .5, .1, .001)$ with $N = 100000$ (see Table 2).

TABLE 2: Baseline disease parameters.

Parameter	Disease 1 Value	Disease 2 Value
N	50000	100000
$1/\mu$	1000 wk	1000 wk
$1/\gamma$	4 wk	2 wk
$1/\sigma$	10 wk	10 wk
β	3.7-wk^{-1} (192-yr^{-1})	14.7-wk^{-1} (764-yr^{-1})

until one goes off in which case a transition is determined by the associated event function.

Figure 2 shows the output for the infected cases of stochastic open SIR. In this section the vector \mathbf{x} represents the sample path for which one has complete information. Assume complete information on the timing and occurrence over a recorded interval of the time series for each individual event propagating through the population. Let $\hat{k} = \sum_{j=1}^{\epsilon} \hat{k}_j$ be the total number of counted events of type \mathcal{P}_j over $[0, T]$. Bookkeep the time and type of event as the set of ordered pairs (t_i, ϵ_i) , where $i = 1, \dots, \hat{k}$, with the t_i in increasing order. Next, consider a recorded event occurring in the ordered interval $[t_{i-1} = t + \tau, t_i = t + \tau + \Delta\tau)$, which was a pathway of type \mathcal{P}_i . In the appendix (Section 1.1) in the supplementary Material available online at doi: 10.1155/2012/390694, it is shown how construction of the likelihood function follows from the stochastic simulation algorithm (SSA) using a factored joint density for any e_{ϵ_i} event tagged with index, ϵ_i , where i is an element of the set consisting of $1, \dots, \epsilon$.

It can also be shown using a factored form of the event function that one can sum over all transitions in the

jump chain resulting from the Kolmogorov forward equation (KFE; see the appendix (Section 1.2)) to obtain,

$$\begin{aligned} \mathcal{L}(\boldsymbol{\alpha}, \mathbf{x}) &= \left\{ \prod_{i=1}^{\hat{k}} \alpha_{\epsilon_i} e_{\epsilon_i}(\mathbf{x}(t_{i-1})) \right\} \times \int_0^T \exp \left\{ -\sum_{i=0}^{\epsilon} \alpha_i e_i dt \right\}, \\ &\propto \prod_{i=1}^{\epsilon} \alpha_i^{\hat{k}_i} \times \exp \left\{ -\alpha_i \int_0^T e_i(\mathbf{x}(t)) dt \right\}, \end{aligned} \quad (6)$$

where ϵ , ϵ_i , \hat{k} , and \hat{k}_i are as defined from above. As shown in the appendix (Sections 1.1–1.3) the standard theory of statistical inference for Markov chains [2, 16, 17] can be applied to simulated Markov processes, to obtain a straightforward, but computationally intensive, maximum likelihood theory for this class of stochastic processes. In fact, these results demonstrate that one can analytically compute closed-form solutions for parameter estimates, since it factors into ϵ independent functions, one for each parameter of an event function and its associated pathway. This gives maximum likelihood estimates of each α_i of the SIR as $\hat{\alpha}_i = k_i / \int_0^T e_i(\mathbf{x}(t)) dt$, for $i = 1, \dots, \epsilon$. This has been demonstrated previously for closed stochastic epidemic models [18, 19]. In this section it has now been shown that similar results hold for *open stochastic endemic disease dynamics*. The factorization will also be utilized in a new way, in a Bayesian context recently advocated by Wilkinson and colleagues [2, 20]. In this case the factorization means that if independent prior distributions are adopted for the parameters this independence will be retained *a posteriori*. Thus, the Bayes theorem may be placed on top of the factorization of likelihood allowing construction of a simulation-based MCMC algorithm for the stochastic SIR. Such an application of this theorem in the SIR case study gives $\alpha_i \mid \mathbf{x} \sim \Gamma\{\alpha_i + k_i, b_i + \int_0^T e_i(\mathbf{x}(t)) dt\}$, where Γ represents the gamma distribution and $i = 1, \dots, 7$ are indexed by each SIR \mathcal{P}_i specified in Table 1 and Figure 1. However, before this method can be applied to the kind of data obtained from actual epidemics the problem of imperfect observation must be addressed. This will be discussed in the next section of this paper.

2.2. SIR Inference: Imperfect Information

2.2.1. Discrete Data Recording. The previous section dealt with the case of availability of perfect information for an observed sample path. In this section the case of imperfect information, such as when sample paths consist of data obtained on fixed recorded intervals, is considered using the output of the vector \mathbf{x} . Thus, the sampled output vector is now considered to contain only partially observed data. A correction can be computed that depends upon the likelihood of a sample path under the true model and the likelihood of the sample path under an approximate model, which takes into account that data are fixed upon two endpoints. This requires computing the likelihood under an inhomogeneous Poisson process model, which will now be stated (Wilkinson [2], Section 10.2).

For simplicity of notation, it is now assumed that the “true” sample path $\mathbf{x}(t)$ is only observed at times $t = 0$ and $t = 1$. Thus, the data fixed upon two endpoints may now be denoted as $\mathbf{x}(0)$ and $\mathbf{x}(1)$. The complete data likelihood for a discretely sampled trajectory on the interval $[0, 1]$ is then approximately given by

$$\begin{aligned} \mathcal{L}^*(\alpha, \mathbf{x}) &= \left(\prod_{i=1}^k \lambda_{\epsilon_i}(t_i) \right) \exp \left\{ -\frac{1}{2} (e_0(\mathbf{x}(0), \alpha) + e_0(\mathbf{x}(1), \alpha)) \right\}, \end{aligned} \quad (7)$$

where $\lambda_j(t) = (1 - t)e_j(\mathbf{x}(0), \alpha) + te_j(\mathbf{x}(1), \alpha)$, $j = 1, \dots, \epsilon$, and represents the rate of the inhomogeneous Poisson process across the interval.

Using the ratio of likelihoods, $\mathcal{L}/\mathcal{L}^*$, allows one to make a robust statistical decision with respect to accepting or rejecting a discretely sampled time interval.

Using a Poisson approximation allows implementing a very fast stochastic simulation algorithm simply (much faster than the standard SSA) by applying probability functions to deterministic flow rates. This essentially corresponds to computing Euler increments for the τ -leap stochastic simulation method [21]. These computational algorithms are briefly described in the appendix (Sections 1.3–1.5).

MCMC implementation using this framework is reasonably straightforward (see [2], Section 10.3): (a) initialize the algorithm with a valid sample path consistent with the observed data. (b) Sample the SIR parameters from their full conditionals given their current sample paths. (c) For each of the reported time intervals propose new sample paths consistent with the reported endpoints and accept/reject it with the Metropolis-Hastings step. (d) Output MCMC state. Go back to (a). Details of the application of this algorithm to the Markovian SIR are discussed in the appendix (Section 1.6).

2.2.2. Nonobservance of Susceptible Cases. Because numbers of susceptible cases are not available from direct observation they must be reconstructed from the epidemic data. For both the simulation and empirical estimation studies a simple reconstruction method [22] is used. This method utilizes the relationship

$$S_{t+\tau} = S_t - C_{t,\tau} + \bar{C}_{t,\tau}, \quad (8)$$

$$S_0 = 0, \quad (9)$$

where S_t is the number of individuals in the susceptible class, $C_{t,\tau}$ the number of reported cases, and $\bar{C}_{t,\tau}$ the average reported number of cases over the entire data set. Given the case report data the susceptible cases are reconstructed by integrating (8) forward from $t = 0$.

3. Results and Discussion

3.1. Reconstructing Stochastic Oscillations. It has long been a challenge in mathematical epidemiology to understand

the recurrence of epidemic outbreaks and establish an appropriate model that allows studying this phenomenon [23–27]. Recurrent epidemics often exhibit intricate and complex dynamics that cannot easily be studied using deterministic models; demographic stochasticity may play a critical role in determining the outcome of the process especially when the population falls below a certain critical size (the critical community size) [28, 29]. Many recent theoretical studies expanding upon Bartlett’s concept of “intrinsic stochastic oscillations” have assumed that the population persists in a long-term stochastic epidemic state [10, 11, 30]; a similar assumption is made in theoretical studies of complex stochastic oscillations in predator-prey systems [31]. This paper will now explore this scenario and estimate parameters for persistent noisy recurrent epidemics using the data assimilation models described in the previous section.

Parameters were estimated for a time series simulated using the stochastic SIR immigration model described previously in this paper. The parameter vector used for α is shown in Table 2 and is representative of a recurrent childhood disease such as a measles, mumps, or pertussis. Two recurrent epidemic scenarios are explored. These are labeled Disease 1 and Disease 2 in Table 2. City sizes of $N = 50000$ and $N = 100,000$ are assumed along with a life expectancy on the order of 20 years. To model the recurrent nature of such an epidemic, an infective immigration rate of $\sigma = .1$ was assumed, so that there is, on average one new infective arriving every 10 weeks. The numbers of infected cases and susceptible cases are always plotted at weekly intervals in the figures. Likewise the sampling interval used to estimate parameters was always made at weekly intervals. This corresponds to the imperfect observation scenario described in Section 2.2.1. In the next section the scenario in which case reports must be used to reconstruct the susceptible class will be dealt with. An example of a simulated infected cases time series is shown in Figure 2: the blue line is from a simulation using Disease 1 parameter values; the red line is from a simulation using Disease 2 parameter values. Two hundred and fifty weeks of observations of infected cases for the SIR immigration model discussed in this paper are shown. Susceptible cases exhibit a similar complex pattern of noisy oscillations but are not shown in the figure.

Using both infected and susceptible cases time series obtained from the simulations the parameters shown in Table 3 were inferred. Analysis of the MCMC data was accomplished using standard Bayesian data analysis [2, 20, 32–34]. Posterior averages and their standard deviations were used to infer parameters after two million MCMC iterations of the inference algorithm. A burnoff of 100000 iterations was made and iterations were thinned every 100 values.

Figure 4 shows Markov chain traces for five hundred weeks of observations of Disease 2. Rapid convergence of the chain towards a region including the target parameters is seen for all SIR parameters (β, γ, μ) except for the immigration rate σ . The color panel shows how estimation of σ improves as data are added (in the figure a yellow line is used to indicate the results for five hundred weeks of observations, a red line for one thousand weeks of observations, and a

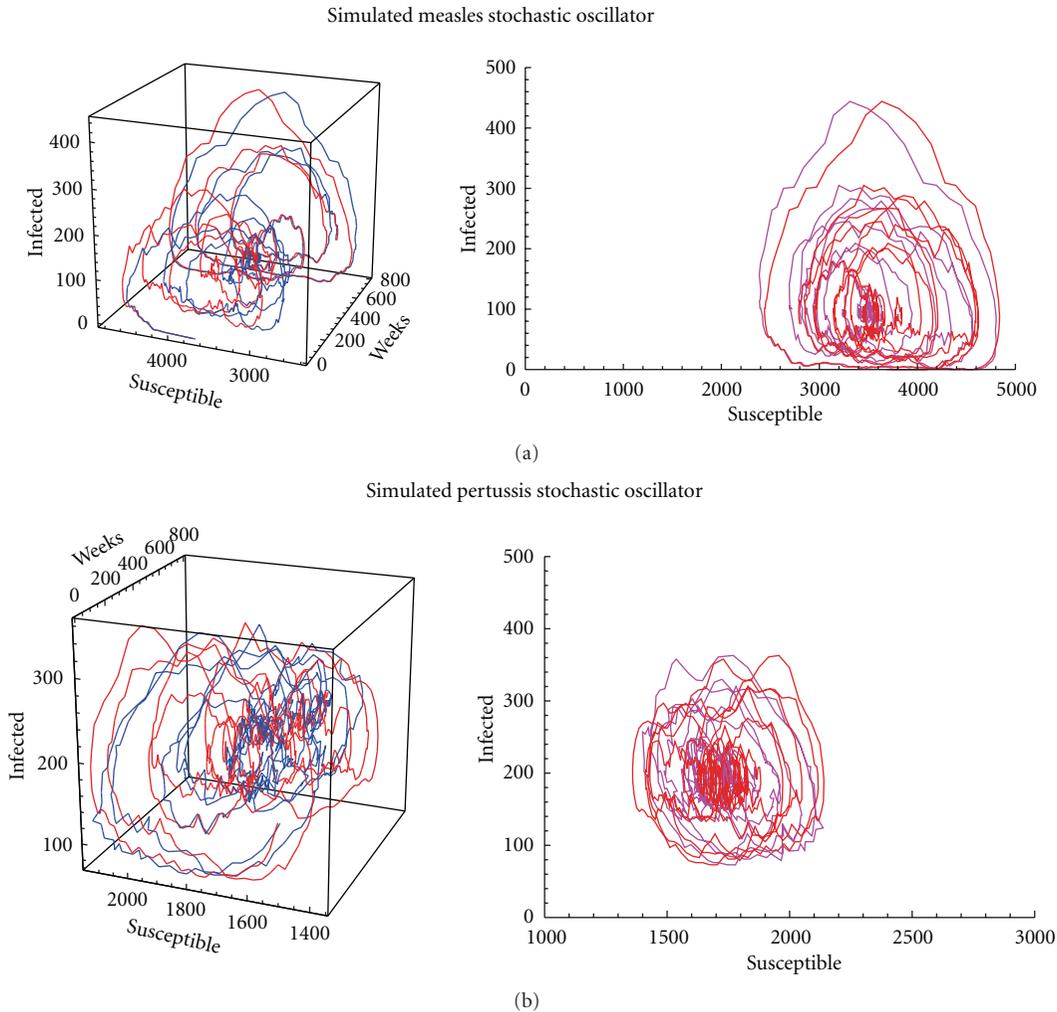


FIGURE 3: Simulation of measles and pertussis as stochastic oscillators with comparison between exact (known) susceptible time series and reconstructed susceptible time series. Parameters used in the simulations are given in Table 2, with pertussis labeled Disease 1 and measles Disease 2. The abbreviation s.o. stands for “stochastic oscillator.” (a) measles stochastic oscillator; (b) pertussis stochastic oscillator.

blue line for ten thousand weeks of observations). Figure 5 shows kernel density estimates for five hundred weeks of observations of Disease 2. The kernel density estimate for migration rate, σ , is for 10000 weeks of observations. Similar results were obtained for Disease 1 type recurrent epidemics (results not shown). All parameters of the epidemic could be estimated for sufficiently long time series (see Table 3).

Finally, it should be pointed out that nearly unbiased estimation of the SIR parameters (β, γ, μ) is sufficient for attractor reconstruction of a persistent recurrent epidemic, at least if the dominant eigenvalue of the point attractor is to be inferred, which is thought to be an important component in driving noisy oscillations of recurrent epidemics in work going back to Soper [23–29].

3.2. Epidemic Inference for 60 UK Cities. In this section parameters are estimated using time series data for 60 UK cities. Pertussis and measles data were obtained using case notification records from the UK Registrar General for England and Wales. Pertussis cases were reported weekly and

biweekly for measles. For both diseases cases reported from the period 1944–1967 were analyzed. City sizes ranged from 10530 (Teignmouth) to 3249440 (London). Reported cases for three UK cities are shown in Figure 7.

Reconstructed susceptible cases (based upon the method described in Section 2.2.2) using simulated measles and pertussis infected time series are shown in Figure 6. Figure 3 shows results from simulation of measles and pertussis stochastic oscillators. Application of this method to perform attractor reconstruction for observed measles time series are shown in Figure 7 for four UK cities. Reasonable similarities were obtained in the comparison between exact (known) susceptible time series versus reconstructed susceptible time series. Parameters used in the simulations are given in Table 2 with pertussis labeled as Disease 1 and measles as Disease 2.

Figure 8 shows the estimates obtained from 60 UK cities for pertussis and measles. Most notable is the large amount of statistical variation seen in the pertussis estimates, particularly in estimates of the duration of infection.

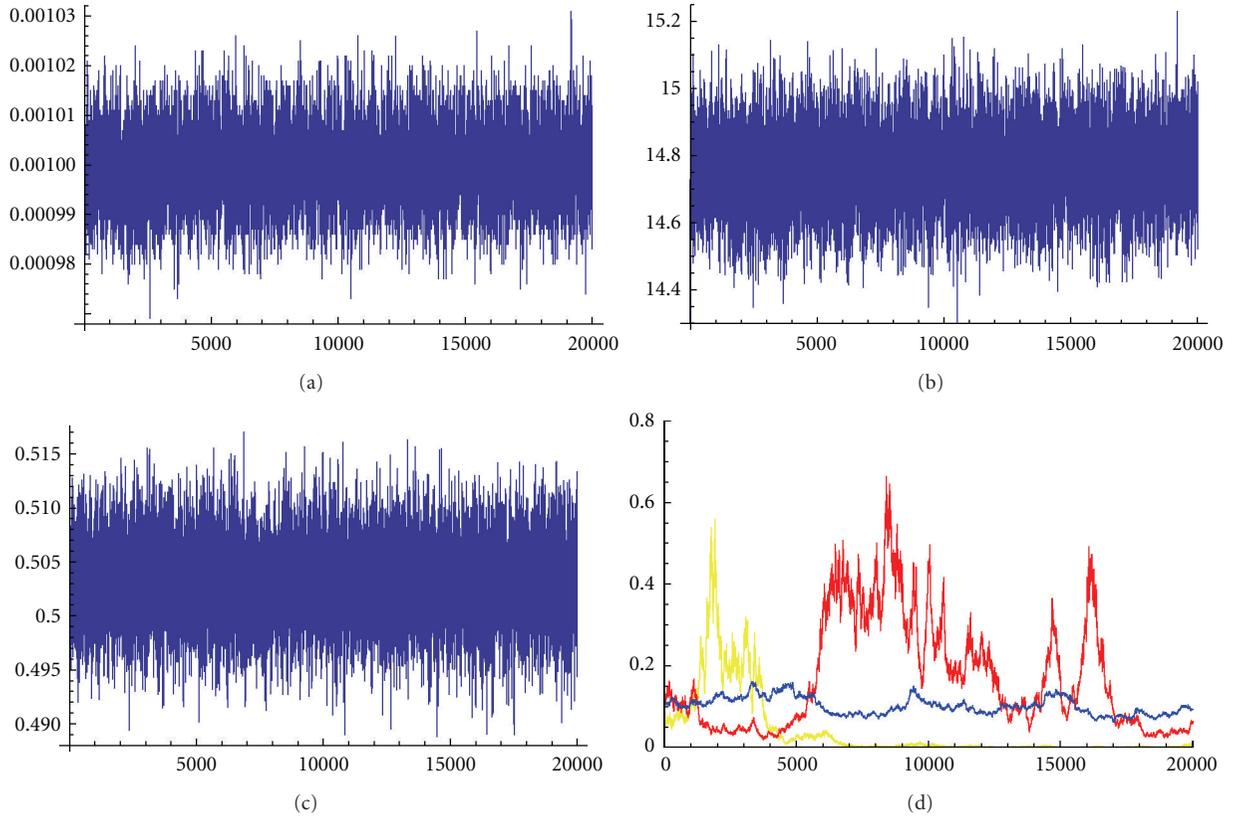


FIGURE 4: Markov chain traces for 500 weeks of observations of Disease 2. The color panel shows how estimation of the migration rate, σ , improves as data are added (yellow line for 500 weeks of observations, red line for 1000 weeks of observations, and blue line for 10,000 weeks of observations). (a) Trace of transmission rate (β); (b) trace of birth rate (N); (c) trace of infection rate (γ); (d) trace of migration rate (σ).

TABLE 3: Posterior estimates of stochastic SIR model.

10000 weeks of observations—Disease 1			
Target	Disease value	Mean	Standard deviation posterior
β	3.70	3.69	.042
γ	.25	.250	.001
σ	.10	.107	.021
μ	.001	.001000	.000004
1000 weeks of observations—Disease 2			
Target	Disease value	Mean	Standard deviation posterior
β	14.7	14.70	.015
γ	.5	.50	.004
σ	.10	.170	.136
μ	.001	.00100	7.18×10^{-6}
10000 weeks of observations—Disease 2			
Target	Disease value	Mean	Standard deviation posterior
β	14.7	14.72	.20
γ	.5	.50	.007
σ	.10	.11	.032
μ	.001	.00100	1.47×10^{-5}

3.3. *Inferring Coherence Resonance.* Coherence resonance occurs when noise is amplified in an otherwise quiescent

system by interaction of the underlying stochasticity of the dynamics with the oscillatory transients of the deterministic dynamics. What has been lacking thus far is a rigorous statistical approach that allows quantifying the theoretical expectations that drive this process using observed time series data. The method developed in this paper is now used to infer endemic sustained oscillations for noisy measles and pertussis epidemics via the mechanism of coherence resonance.

Kuske et al. [11] showed that the Poisson process model of the SIR may be approximated using a stochastic ordinary differential equation with a change of variables. The linearization of scaled model has oscillations that are slowly decaying with unity frequency. Kuske et al. conjecture that solutions of the stochastic SIR model resemble a different approximate model which captures the essence of the full stochastic model. In this stochastic analogue the sustained oscillations have a very particular structure: they are a family of sinusoids modulated by the Ornstein-Uhlenbeck processes. Making this conjecture Kuske et al. [11] derive simple quantitative conditions for the existence of sustained oscillations in noisy time series. Hence, they are able to describe the parameter region for R_0 and γ in detail, including the behavior of the power spectral density of stochastic model and its multiscale approximation. Their parameter space will be taken as the starting point

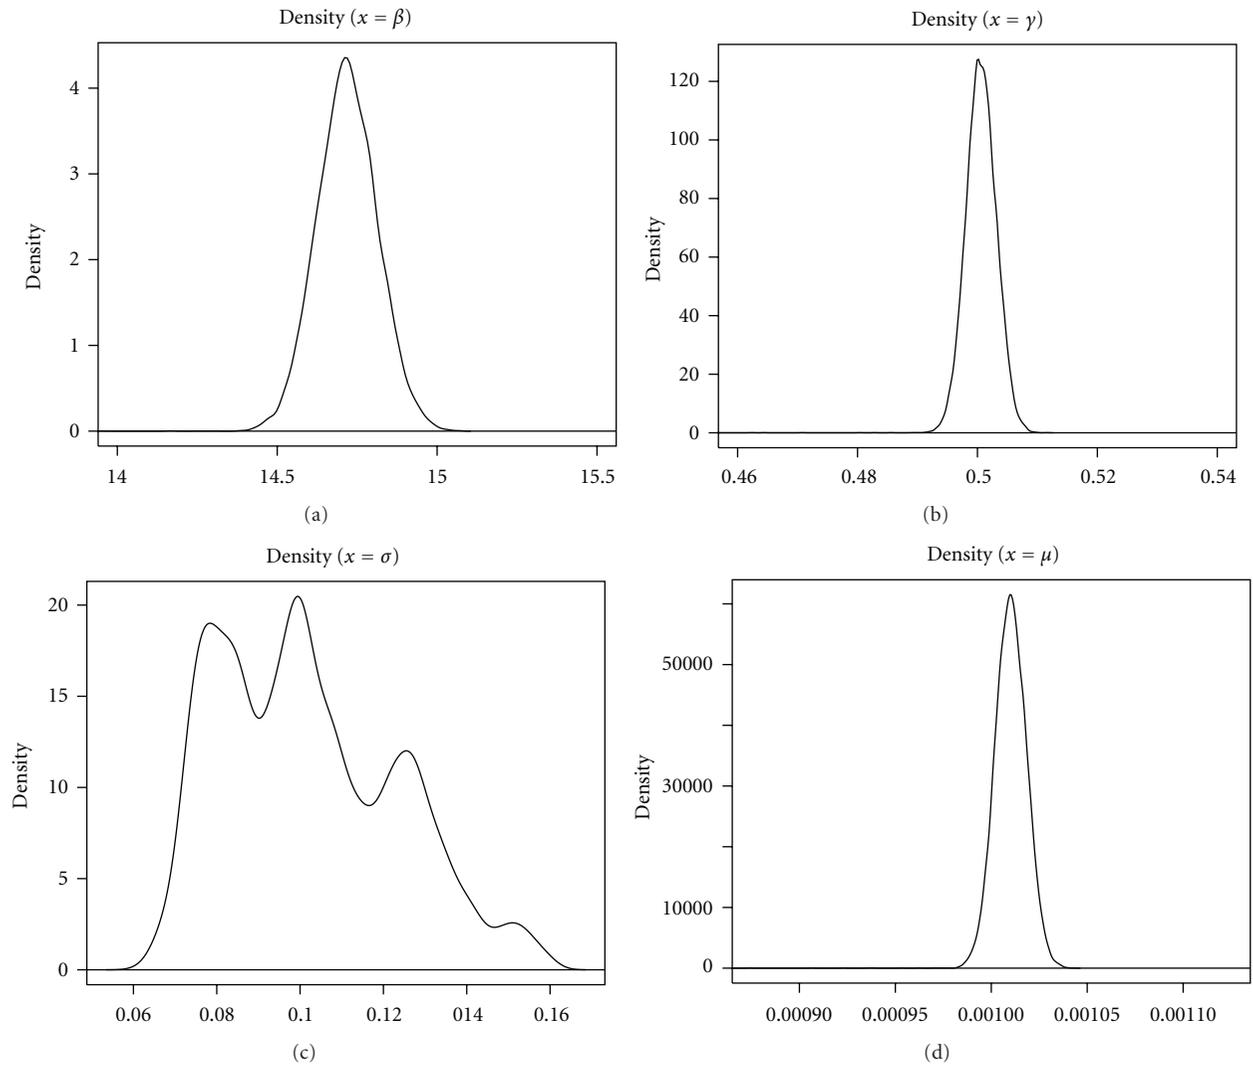


FIGURE 5: Kernel density estimates for 500 weeks of observations of Disease 2. The kernel density estimate for migration rate, σ , is for 10,000 thousand weeks of observations. (b) Density of of transmission rate (β); (a) density of infection rate (γ); (c) density of migration rate (σ); (d) density of birth rate (μ).

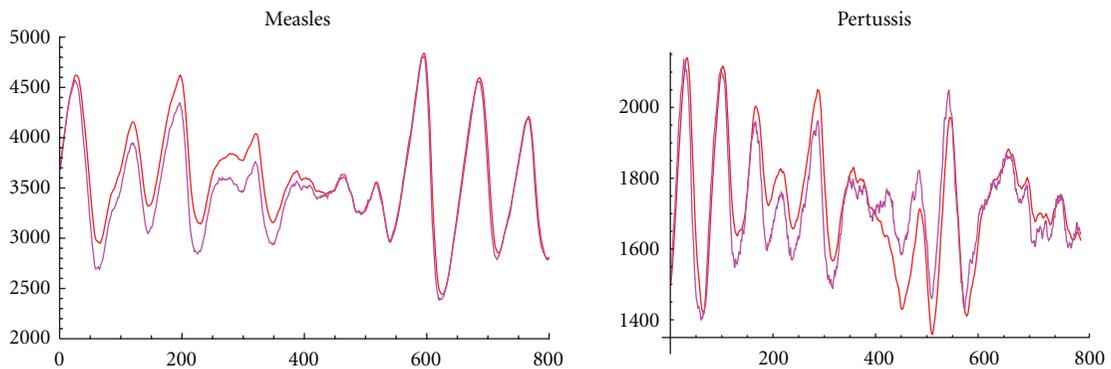
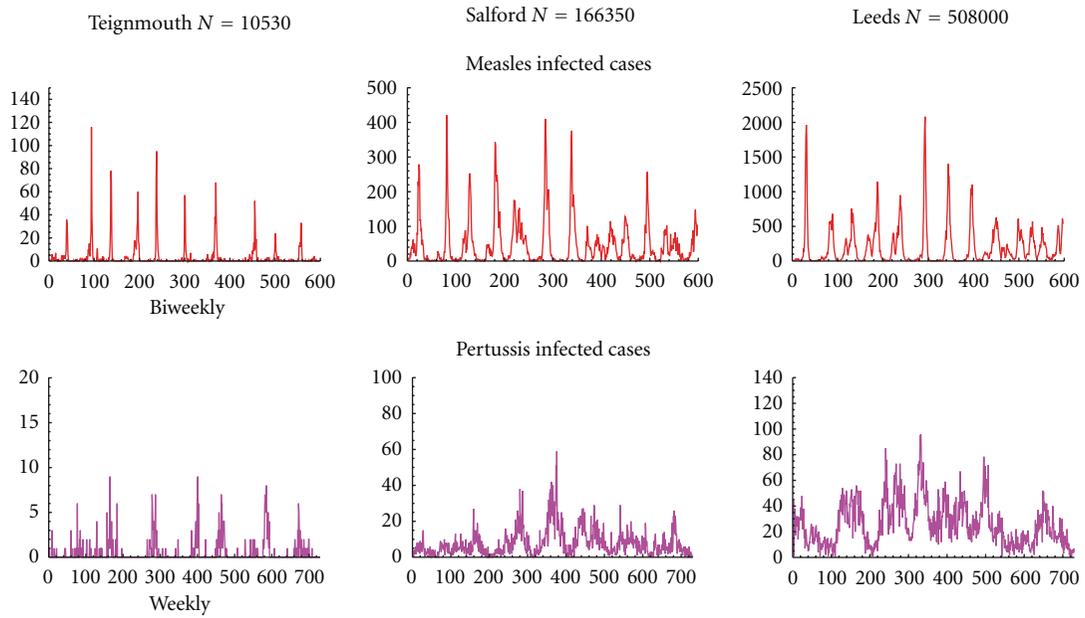
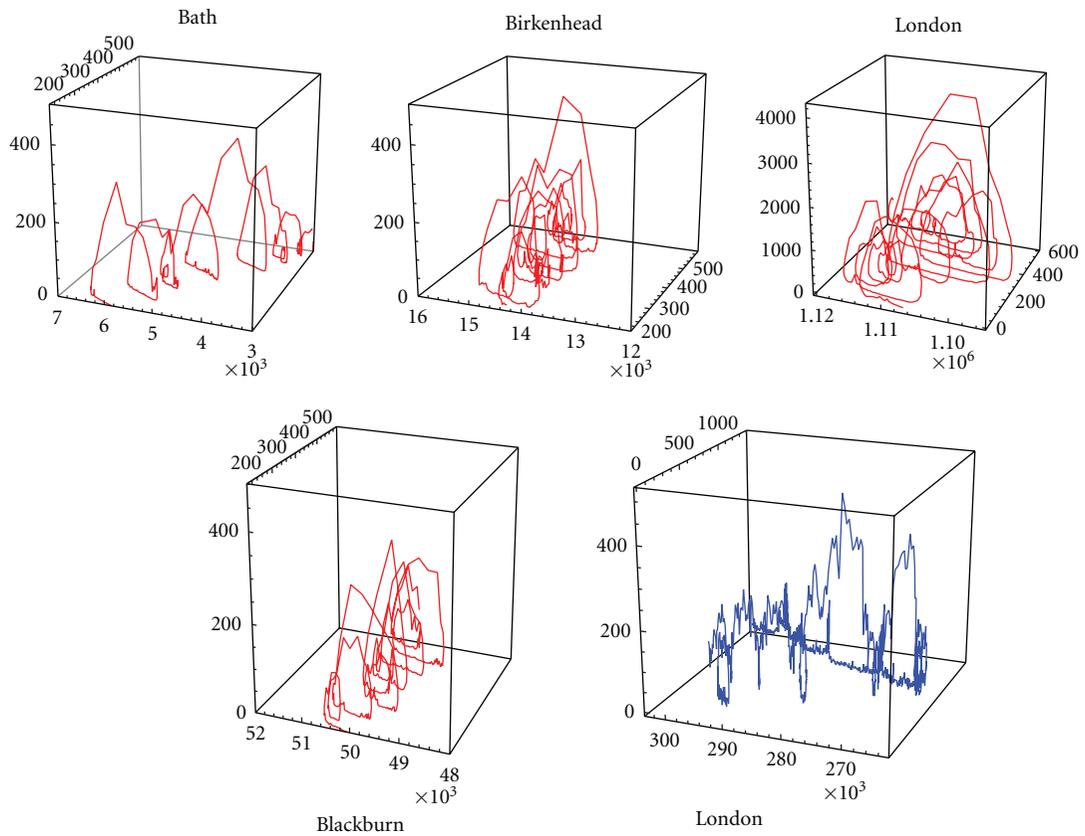


FIGURE 6: Measles and pertussis reconstructed susceptible time series. Cases are on x -axis, weeks are plotted on the x -axis. Blue: susceptible time series. Red: reconstructed susceptible time series.



(a)



(b)

FIGURE 7: (a) Measles and pertussis cases (y -axis) plotted as time series in three UK cities. Measles infected cases were reported as biweekly cases, while pertussis cases were reported weekly. (b) Measles (red) time series as reconstructed stochastic oscillators for four UK cities. The reconstructed oscillator for pertussis is shown in blue for London.

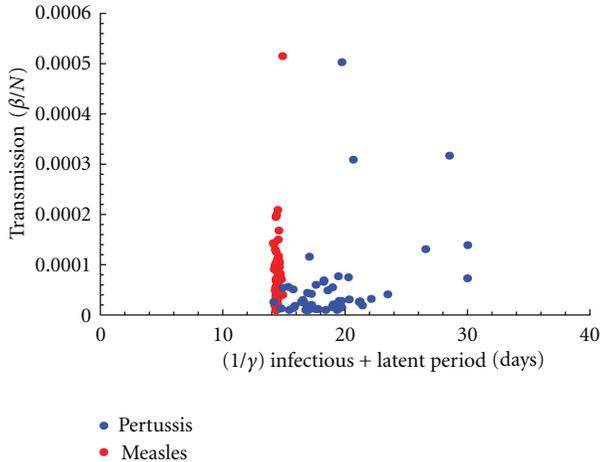


FIGURE 8: Measles and pertussis estimates of transmission (β) and infectious period ($1/\gamma$) epidemic parameters for 60 UK cities.

for formulating the hypothesis of coherence resonance in stochastic epidemics explored in this paper.

Kuske et al. [11] give the biological criterion for sustained oscillations via coherence resonance for the SIR model in terms of two bounds:

$$\epsilon^2 = \frac{R_0}{2} \sqrt{\frac{\mu}{\mu + \gamma} \frac{1}{R_0 - 1}} \ll 1, \quad (10)$$

$$\frac{\delta^2}{2\epsilon^2} = \frac{\mu + \gamma}{4N\mu} \left(1 + \frac{R_0 + 1}{R_0 - 1} + 2 \frac{\mu + \gamma}{\mu(R_0 - 1)} \right) \ll 1.$$

Hence, these bounds can be explored by estimating R_0 and γ that explore this region for the UK measles and pertussis data. Assume that $1/\mu = 70$ is a nuisance parameter. Although σ was estimated, it does not play a role in the following analysis; therefore it is ignored in this section. In addition to μ parameter estimates of interest are those of γ , β , and R_0 which are required to estimate ϵ^2 and $\delta^2/2\epsilon^2$. The major predictions with respect to stochastic amplification in the model of Kuske et al. [11, Page 465] are as follows: (1) for very small values of $\delta^2/2\epsilon^2$ one expects to see very small oscillations. (2) When $\delta^2/2\epsilon^2$ is increased but below one the stochastic fluctuations balance with the deterministic slow decay so that both stochastic and deterministic components interplay to determine the attractor dynamics. (3) When $\delta^2/2\epsilon^2$ is large the stochastic variations govern the dynamics so that an approximation based upon slowly varying modulations is no longer appropriate.

The key results are as follows. Figure 9 shows the estimated variance of the stationary process from measles and pertussis time series. Since this quantity restricts the variance of the stochastic fluctuations relative to the slow time scale it can be used to determine the relative sensitivity of fluctuations on this time scale. It may be observed in Figure 9 that there exist very small estimated values $\delta^2/2\epsilon^2$ for pertussis (blue); hence, one expects to see relatively small very noisy oscillations propagated through the attractor. In this case the demographic noise will not be likely to be amplified optimally with respect to the deterministic

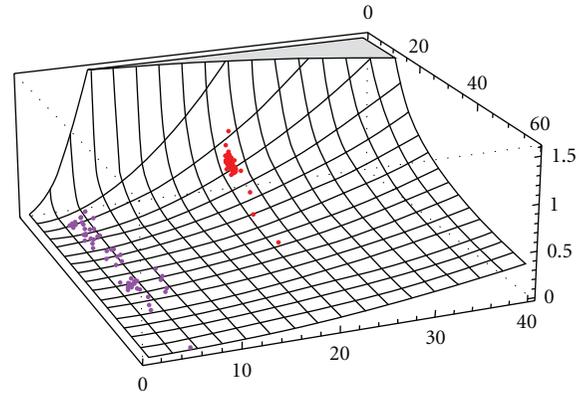


FIGURE 9: Estimated variance of the stationary process from measles and pertussis time series. Plotted on the x -axis is the rate of infection (γ) in years. Plotted on the z -axis is the variance ($\delta^2/2\epsilon^2$). Plotted y -axis is the reproductive rate of the disease (R_0). One observes very small estimated values $\delta^2/2\epsilon^2$ for pertussis (blue), hence one expects to see relatively small very noisy oscillations propagated through the attractor. In contrast, measles (red) exhibits more moderate estimated values $\delta^2/2\epsilon^2$. This implies that the stochastic fluctuations balance with the deterministic slow decay so that both stochastic and deterministic processes contribute the dynamics in terms of producing patterns of coherence resonance.

frequency in the power spectrum and will show more irregular fluctuations due to stochastic amplification of demographic noise. Hence, it may be predicted that the power spectral distribution will not be as sharply peaked and that the multiscale approximation is not as valid for pertussis as for other pathogens. In contrast, Figure 10 shows that for measles (red) one observes more moderate estimated values of $\delta^2/2\epsilon^2$. This implies that the stochastic fluctuations balance with the deterministic slow decay so that both stochastic and deterministic processes contribute the dynamics in terms of producing patterns of coherence resonance. For measles epidemics it is predicted that the power spectral density will have stronger peaks in the vicinity of the deterministic frequency. Measles noisy oscillations are predicted to be better structured and exhibit more coherent cycles around the endogenous period and measles epidemics will exhibit more sensitivity to stochastic amplification. That is they will amplify the noise to generate more regular stochastic cycles in the neighborhood of a fixed frequency.

Figure 10 shows a plot of estimated per year rate of infection (labeled as gamma) versus R_0 in analytically predicted bounds expected for multiscale dynamics leading to coherence resonance. In Figure 9 the light green line represents the V-shaped boundary of $\epsilon^2 < .1$ computed using $N = 500000$ and $\mu = 1/55$. This region is approximately the same size when computed for ranges of N between 500000 and 2,000000 [11]. The blue line in Figure 7 represents the contour of the bound $\delta^2/2\epsilon^2 < .2$. Both measles (red) and pertussis (blue) estimates lie well within the bound set by $\delta^2/2\epsilon^2$ for coherence resonance; however, pertussis lies on the boundary of the ϵ^2 bound, which seems to suggest that these epidemics are *not* as likely to exhibit

multiscale dynamics as measles epidemics. There does not appear to exist quite a strong separation between slow and fast time scales in determining pertussis dynamics as there does for measles dynamics. Hence, one expects less coherence and less structured oscillations for pertussis more coherence and structured oscillation for measles epidemics. These results are supported by those observed in Figure 10 and complement each other.

3.4. Systems Biology Approaches to Inference. This paper utilizes a parameter estimation used for mechanistic modeling of biochemical systems [2] to address the important challenge of bridging the gap that exists between mathematical modeling of epidemics and data analysis. In this paper the Bayesian MCMC method has been shown to be useful in bridging such a gap as well as in testing interesting hypotheses regarding the properties of stochastic amplification in epidemics. However, the application of this computational theory in this paper to a simplified open Markovian SIR is really just a first step. But it is an important one and has allowed investigating the properties of data from endemic diseases—a highly nontrivial inference problem in epidemiology. In this section some other, more recent systems biology methods are reviewed and compared to the method used in this paper. Some advantages of systems biology inference methods will be briefly discussed and may be used building upon the results presented in this paper.

The application of computational and mathematical techniques from what has been called algorithmic systems biology [35] to an epidemiological modeling problem will likely prove fruitful. The derivations of stochastic modeling of continuous time processes and the corresponding likelihoods are quite general. However, the approach by Wilkinson [2] and colleagues was a first step in modeling systems in which stochastic effects due to small numbers of molecules or individuals in populations are to be studied. In fact, subsequent studies by the authors focused on inference methods based upon diffusion approximations, which are more tractable and scale up to large systems more easily but are not appropriate for systems in which low densities are common. Applying the Wilkinson Bayesian MCMC Markov jump process approach requires approximating a continuous system using a discrete Poisson approximation. However, as shown in this paper such an approximation does allow obtaining results for endemic diseases which would otherwise be impossible to obtain using earlier algorithms put forward such as that by Gibson and Renshaw [9] for example. Also even as the smaller scale used in this paper it is so computationally intensive it cannot yet be applied to larger scale problems. That is the main reason in this paper that a simplified Markovian SIR model was used. However, even for simplified models one may have the problem of taking long waiting times for rare events. Both simulated and variational maximum likelihood methods in systems biology [36–38] suffer from similar maladies at the Bayesian MCMC methods.

Recent breakthroughs in automated estimation of rare event probabilities in biochemical systems [39–41], however, may allow addressing some of these fundamental problems.

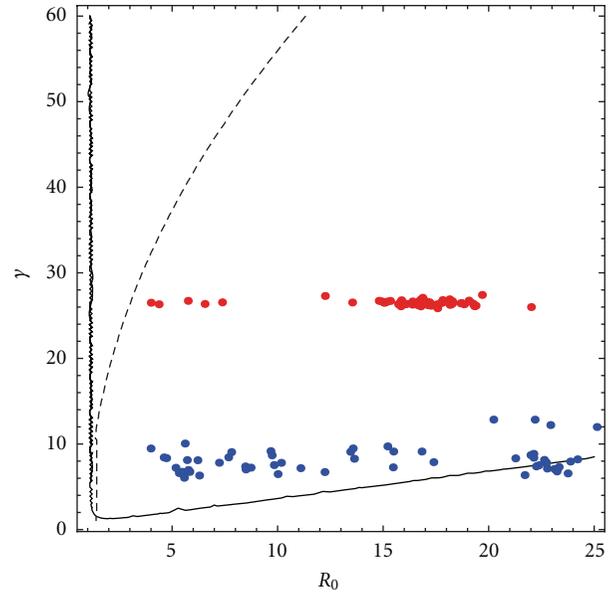


FIGURE 10: Plot of estimated per year rate of infection (labeled gamma on the y -axis) versus R_0 (labeled on the x -axis) in analytically predicted bounds expected for multiscale dynamics leading to coherence resonance. The light green line represents the V-shaped boundary of $\epsilon^2 < .1$ computed using $N = 500000$ and $\mu = 1/55$. This region is approximately the same size when computed for ranges of N between 500000 and 2,000000 [11]. The blue line represents the contour of the bound $\delta^2/2\epsilon^2 < .2$. Both measles (red) and pertussis (blue) estimates lie well within the bound set by $\delta^2/2\epsilon^2$ for coherence resonance; however, pertussis lies on the boundary of the ϵ^2 bound, which seems to suggest that these epidemics are *not* as likely to exhibit multiscale dynamics as measles epidemics. This can be seen visually by comparing the reconstruction of the stochastic oscillators for measles (red) and pertussis (blue), respectively, for London shown in Figure 7.

For the first time an accelerated maximum likelihood estimation for stochastic biochemical systems is in sight that can be based on the continuous time SSA. Construction of inference algorithms based upon these recent studies in systems biology will allow extending the results presented in this paper to more realistic models of the epidemiological process such as including multiple exposed and infected classes. It will also allow including the possibility of disease interactions which, for two diseases can require up to fifty state variables to model [42].

4. Conclusion

In this paper a straightforward Bayesian MCMC methodology for inferring parameters for open SIR models using stochastic simulation is applied to both simulated and observed epidemic time series data. The methods described in this paper are general enough for extension to more complex epidemiological scenarios, which is currently the goal of future work. This is useful because the efficient integration of complex likelihoods for population models is currently an object of intense ongoing research. Analysis

of the data for the methodology developed in this paper is accomplished using standard Bayesian data analysis [2, 20, 32–34].

The results obtained in this paper show how pertussis and measles epidemics behave with respect to the presence of demographic noise. Time series for 60 UK cities were used to estimate epidemiological parameters for these pathogens. A coherence resonance model was fit to the data to infer the role of multiscale effects in producing period and amplitude in the epidemics. It was found that measles appears to fit the model rather well. However, pertussis does not seem to fit the model, and it is predicted that there does not appear to exist quite a strong separation between slow and fast time scales as for pertussis as seems to exist for measles epidemics. Therefore, one expects less coherence and less structured oscillations for pertussis but more coherence and structured oscillation for measles epidemics. The statistical theory developed in this paper was used to investigate coherence resonance of epidemics [10, 11] using empirical time series data. It is hoped that future work will be directed toward extending these results to more complex epidemic modeling [43] such as theory of immune-mediated processes in pathogen interactions [42, 44].

Acknowledgments

This work was partly funded by the University of Missouri and Duke University. The author would like to thank Helen Wearing and Pej Rohani for comments and suggestions on earlier versions of this paper.

References

- [1] J. S. Clark, *Models for Ecological Data*, Princeton University Press, Princeton, NJ, USA, 2007.
- [2] D. J. Wilkinson, *Stochastic Modelling for Systems Biology*, Chapman and Hall, Boca Raton, Fla, USA, 2006.
- [3] H. Andersson and T. Britton, *Stochastic Epidemic Models and Their Analysis*, Springer, Berlin, Germany, 2000.
- [4] P. D. O'Neill and G. O. Roberts, "Bayesian inference for partially observed stochastic epidemics," *Journal of the Royal Statistical Society. Series A*, vol. 162, no. 1, pp. 121–129, 1999.
- [5] P. D. O'Neill, "A tutorial introduction to Bayesian inference for stochastic epidemic models using Markov chain Monte Carlo methods," *Mathematical Biosciences*, vol. 180, pp. 103–114, 2002.
- [6] N. G. Becker, *Analysis of Infectious Disease Data*, Chapman and Hall, London, UK, 1989.
- [7] N. G. Becker and T. Britton, "Statistical studies of infectious disease incidence," *Journal of the Royal Statistical Society. Series B*, vol. 61, no. 2, pp. 287–307, 1999.
- [8] S. Cauchemez and N. M. Ferguson, "Likelihood-based estimation of continuous-time epidemic models from time-series data: application to measles transmission in London," *Journal of the Royal Society Interface*, vol. 5, no. 25, pp. 885–897, 2008.
- [9] G. J. Gibson and E. Renshaw, "Estimating parameters in stochastic models using Markov chain methods," *IMA J. Math. Appl. Med. Biol.*, vol. 15, pp. 19–40, 1998.
- [10] D. Alonso, A. J. McKane, and M. Pascual, "Stochastic amplification in epidemics," *Journal of the Royal Society Interface*, vol. 4, no. 14, pp. 575–582, 2007.
- [11] R. Kuske, L. F. Gordillo, and P. Greenwood, "Sustained oscillations via coherence resonance in SIR," *Journal of Theoretical Biology*, vol. 245, no. 3, pp. 459–469, 2007.
- [12] W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," *Proceedings of the Royal Society of London Series A*, vol. 115, pp. 700–721, 1927.
- [13] R. M. Anderson and R. M. May, *Infectious Diseases of Humans: Dynamics and Control*, Oxford University Press, Oxford, UK, 1991.
- [14] A. G. McKendrick, "Applications of mathematics to medical problems," *Proceedings of the Edinburgh Mathematical Society*, vol. 44, pp. 98–130, 1926.
- [15] M. S. Bartlett, "Some evolutionary stochastic processes," *Journal of the Royal Statistical Society: Series B*, vol. 11, pp. 211–229, 1949.
- [16] P. Billingsley, *Statistical Inference for Markov Processes*, University of Chicago, Chicago, Ill, USA, 1961.
- [17] P. Guttorp, *Stochastic Modeling of Scientific Data*, Chapman and Hall, Boca Rotan, Fla, USA, 1995.
- [18] W. N. Rida, "Asymptotic properties of some estimators for the infection rate in the general stochastic epidemic model," *Journal of the Royal Statistical Society: Series B*, vol. 53, pp. 269–283, 1991.
- [19] R. J. Kryscio, "The transition probabilities of the general stochastic epidemic model," *Journal of Applied Probability*, vol. 12, pp. 415–424, 1975.
- [20] R. J. Boys, D. J. Wilkinson, and T. B. L. Kirkwood, "Bayesian inference for a discretely observed stochastic kinetic model," *Statistics and Computing*, vol. 18, no. 2, pp. 125–135, 2008.
- [21] D. T. Gillespie, "Approximate accelerated stochastic simulation of chemically reacting systems," *Journal of Chemical Physics*, vol. 115, no. 4, pp. 1716–1733, 2001.
- [22] G. V. Bobashev, S. P. Ellner, D. W. Nychka, and B. T. Grenfell, "Reconstructing susceptible and recruitment dynamics from measles epidemic data," *Mathematical Population Studies*, vol. 8, no. 1, pp. 1–29, 2000.
- [23] W. H. Hamer, "Epidemic disease in England—the evidence of variability and of persistence of type," *Lancet*, vol. 1, pp. 733–739, 1906.
- [24] H. E. Soper, "The interpretation of periodicity in disease prevalence," *Journal of the Royal Statistical Society: Series B*, vol. 92, pp. 34–73, 1929.
- [25] M. S. Bartlett, *Stochastic Models in Ecology and Epidemiology*, Methuen, London, UK, 1960.
- [26] M. S. Bartlett, "Chance or chaos?" *Journal of the Royal Statistical Society: Series A*, vol. 153, pp. 321–349, 1990.
- [27] N. T. J. Bailey, *The Mathematical Theory of Infectious Disease and Its applications*, Griffin, London, UK, 1975.
- [28] I. Nåsell, "On the time to extinction in recurrent epidemics," *Journal of the Royal Statistical Society. Series B*, vol. 61, no. 2, pp. 309–330, 1999.
- [29] I. Nasell, "Measles outbreaks are not chaotic," in *Mathematical Approaches for Emerging and Reemerging Infectious Disease: Models, Methods and Theory*, Springer, 2002.
- [30] N. G. V. Kampen, *Stochastic Processes in Physics and Chemistry*, North-Holland, Amstserdam, The Netherlands, 1981.
- [31] A. J. McKane and T. J. Newman, "Predator-prey cycles from resonant amplification of demographic stochasticity," *Physical Review Letters*, vol. 94, no. 21, Article ID 218102, pp. 1–4, 2005.
- [32] B. P. Carlin and T. A. Louis, *Bayes and Empirical Bayes Methods for Data Analysis*, CRC Press, Boca Raton, Fla, USA, 2000.
- [33] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis*, CRC Press, Boca Raton, Fla, USA, 2004.

- [34] J. S. Clark, "Why environmental scientists are becoming Bayesians," *Ecology Letters*, vol. 8, no. 1, pp. 2–14, 2005.
- [35] C. Priami, "Algorithmic systems biology," *Communications of the ACM*, vol. 52, no. 5, pp. 80–88, 2009.
- [36] S. K. Poovathingal and R. Gunawan, "Global parameter estimation methods for stochastic biochemical systems," *BMC Bioinformatics*, vol. 11, article 414, 2010.
- [37] T. Tian, S. Xu, J. Gao, and K. Burrage, "Simulated maximum likelihood method for estimating kinetic rates in gene expression," *Bioinformatics*, vol. 23, no. 1, pp. 84–91, 2007.
- [38] Y. Wang, S. Christley, E. Mjolsness, and X. Xie, "Parameter inference for discretely observed stochastic kinetic models using stochastic gradient descent," *BMC Systems Biology*, vol. 4, article 99, 2010.
- [39] B. J. Daigle, M. K. Roh, D. T. Gillespie, and L. R. Petzold, "Automated estimation of rare event probabilities in biochemical systems," *Journal of Chemical Physics*, vol. 134, no. 4, Article ID 044110, 2011.
- [40] H. Kuwahara and I. Mura, "An efficient and exact stochastic simulation method to analyze rare events in biochemical systems," *Journal of Chemical Physics*, vol. 129, no. 16, Article ID 165101, 2008.
- [41] D. T. Gillespie, M. Roh, and L. R. Petzold, "Refining the weighted stochastic simulation algorithm," *Journal of Chemical Physics*, vol. 130, no. 17, Article ID 174103, 2009.
- [42] D. A. Vasco, H. J. Wearing, and P. Rohani, "Tracking the dynamics of pathogen interactions: modeling ecological and immune-mediated processes in a two-pathogen single-host system," *Journal of Theoretical Biology*, vol. 245, no. 1, pp. 9–25, 2007.
- [43] M. J. Keeling and P. Rohani, *Modelling Infectious Diseases*, Princeton University Press, Princeton, NJ, USA, 2007.
- [44] P. Rohani, C. J. Green, N. B. Mantilla-Beniers, and B. T. Grenfell, "Ecological interference between fatal diseases," *Nature*, vol. 422, no. 6934, pp. 885–888, 2003.

Research Article

Higuchi Fractal Properties of Onset Epilepsy Electroencephalogram

Truong Quang Dang Khoa,¹ Vo Quang Ha,¹ and Vo Van Toi^{1,2}

¹Biomedical Engineering Department, International University of Vietnam National University, Ho Chi Minh City, Vietnam

²Biomedical Engineering Department, Tufts University, MA 02155, USA

Correspondence should be addressed to Truong Quang Dang Khoa, khoa@ieee.org

Received 27 August 2011; Revised 23 October 2011; Accepted 30 October 2011

Academic Editor: Ranjit Kumar Upadhyay

Copyright © 2012 Truong Quang Dang Khoa et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Epilepsy is a medical term which indicates a common neurological disorder characterized by seizures, because of abnormal neuronal activity. This leads to unconsciousness or even a convulsion. The possible etiologies should be evaluated and treated. Therefore, it is necessary to concentrate not only on finding out efficient treatment methods, but also on developing algorithm to support diagnosis. Currently, there are a number of algorithms, especially nonlinear algorithms. However, those algorithms have some difficulties one of which is the impact of noise on the results. In this paper, in addition to the use of fractal dimension as a principal tool to diagnose epilepsy, the combination between ICA algorithm and averaging filter at the preprocessing step leads to some positive results. The combination which improved the fractal algorithm become robust with noise on EEG signals. As a result, we can see clearly fractal properties in preictal and ictal period so as to epileptic diagnosis.

1. Introduction

Fractal dimension (FD) is considered as an important parameter applied to human biosignals. The results of FD in time domain depend on algorithm and window length. This problem was analyzed deeply by Pradhan and Dutt, when they discussed the effect of window length and window displacement on results [1].

In 2001, Echaz et al. [2] compared between results of Higuchi algorithm [3], Katz [4] and Petrosian algorithm [5] in intracranial electroencephalogram (I-EEG) epilepsy signal. The results showed that Katz's algorithm was the most consistent method for discrimination of epileptic states from the I-EEG, likely due to its exponential transformation of FD values and relative insensitivity to noise. Higuchi's method, however, yields a more accurate estimation of signal FD, when tested on synthetic data, but is more sensitive to noise. Petrosian's method performance depends on the type of binary sequence used. If a binary sequence based on slope-sign changes is utilized then this method becomes less suitable for analog signal analysis, given its high sensitivity to noise and its poor reproducibility of dynamic range of

synthetic FDs. Kannathal et al. [6] used Katz algorithm and Higuchi algorithm to calculate averaging fractal dimensions of 2 groups: one is healthy, another is epilepsy patient. Results show that the FDs of the epilepsy group are lower than healthy one in both methods. In epilepsy detection, Esteller et al. [7] said that by the time seizures happened, the fractal dimension using Katz algorithm increases in the ictal period, followed by a fall to the lowest complexity level of the recording. Moreover, in 2003, this group used 6 parameters, including curve length, energy, nonlinear energy, spectral entropy, sixth power, and energy of wavelets packets, as features for EEG segmentation in epilepsy [8]. In the same way, Bao et al. used Higuchi Algorithm, Petrosian algorithm, Hjorth parameters, power spectra, means, standard deviation, and neural network for epilepsy diagnosis [9].

In this study, we analyze the fractal properties as parameters for both EEG and ECG epilepsy detection.

2. Methodology

In this paper, we proposed two methods to analyze epilepsy data. The first method includes two steps: all of channels were

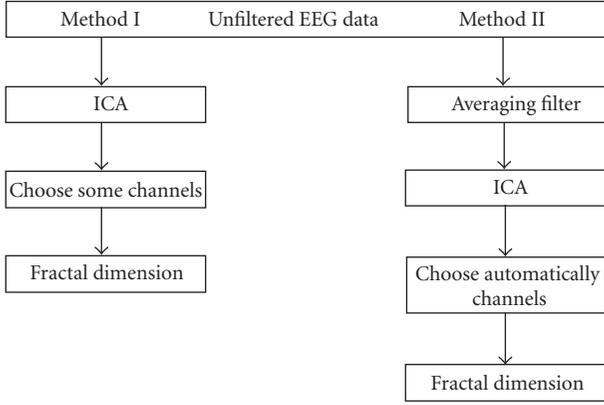


FIGURE 1: Proposed algorithms.

analyzed to archive independent components by ICA algorithm. After that, Higuchi algorithm was used to calculate fractal dimension. The second method processed the same way to the first method, except that an averaging filter was used as the first. The methodology used in this paper consists of the steps shown in the diagram in Figure 1.

2.1. Averaging Filter (AF). The averaging filter is the simplest type of low-pass filter using when the neighborhood considered is too large blurring and other unwanted effects can appear in the data set. This method can be useful to avoid very high frequency noise and white noise. The value of a sample is calculated by the average of its neighbors:

$$x_n = \frac{1}{2k+1} \sum_{i=-k}^k x_{n+i}, \quad (1)$$

where k is the window length, x_n is the value of n th sample.

By experiments, we assume that $k = 3$ is suitable for epilepsy prediction.

2.2. Independent Component Analysis (ICA)

2.2.1. Definition of ICA. We assume that we observe n linear mixtures x_1, \dots, x_n of n independent components:

$$x_j = a_{j1}s_1 + a_{j2} + \dots + a_{jn}s_n, \quad j = \overline{1, n}. \quad (2)$$

We have now dropped the time index t ; in the ICA model, we assume that each mixture x_j as well as each independent component s_k is a random variable, instead of a proper time signal [10]. Without loss of generality, we can assume that both the mixture variables and the independent components have zero mean; if this is not true, then the observable variables x_i can always be centered by subtracting the sample mean, which makes the model zero mean:

$$\tilde{x} = x - E(x). \quad (3)$$

Let x be the random vectors whose elements are the mixtures x_1, \dots, x_n and let s be the random vector with the

components s_1, \dots, s_n . Let A be the matrix containing the elements a_{ij} . The model can now be written as follows:

$$x = As \quad \text{Or} \quad x = \sum_{i=1}^n a_i s_i. \quad (4)$$

The above equation is called independent component analysis or ICA. The problem is to determine both the matrix A and the independent components s , knowing only the measured variables x . The only assumption the methods take is that the components s_i are independent. It has also been proved that the components must have nongaussian distribution.

Before the application of the ICA algorithm (and after centering), we transform the observed vector x linearly to obtain a new vector \tilde{x} which is white (its components are uncorrelated and their variances equal unity).

Whitening can be performed via eigenvalue decomposition of the covariance matrix:

$$E\{xx^T\} = EDE^T, \quad (5)$$

where E is the orthogonal matrix of eigenvectors of $E\{xx^T\}$ and D is the diagonal matrix of its eigenvalues, $D = \text{diag}(d_1, \dots, d_n)$. Whitening can now be done by

$$\tilde{x} = ED^{-1/2}E^T x. \quad (6)$$

2.2.2. Fast ICA for n Units [10]. A unit represents a processing element, for example, an artificial neuron with its weights W .

To estimate several independent components, the weights w_1, \dots, w_n must be determined. The problem is that the outputs $w_1^T x, \dots, w_n^T x$ must be done as independent as possible after each iteration in order to avoid the convergence to the same maxima. One method is to estimate the independent components one by one.

Algorithm.

Step 1. Initialize w_i .

Step 2. Newton phase:

$$w_i = E\{\tilde{x}g(w_i^T \tilde{x})\} - E\{g'(w_i^T \tilde{x})\} w_i, \quad (7)$$

where g is a function with one of the following forms:

$$g_1(y) = \tanh(a_1 y),$$

$$g_2(y) = y \exp\left(-\frac{1}{2}y^2\right), \quad g_3(y) = 4y^3. \quad (8)$$

Step 3. Normalization:

$$w_i = \frac{1}{\|w_i\|} w_i. \quad (9)$$

Step 4. Decorrelation:

$$w_i = w_i - \sum_{j=1}^{i-1} w_i^T w_j w_j. \quad (10)$$

Step 5. Normalization (like in the Step 3).

Step 6. Go to Step 2 if not converged.

2.2.3. Higuchi's Fractal Dimension Algorithm. Higuchi's algorithm calculates fractal dimension of a time series directly in the time domain. It is based on a measure of length, $L(k)$, of the curve that represents the considered time series while using a segment of k samples as a unit, if $L(k)$ scales like

$$L(k) \sim k^{-D_f}. \quad (11)$$

The curve is said to show fractal dimension D_f because a simple curve has dimension equal 1 and a plane has dimension equal 2; value of D_f is always between 1 (for a simple curve) and 2 (for a curve which nearly fills out the whole plane). D_f measures complexity of the curve and so of the time series this curve represents on a graph.

From a given time series, $X(1), X(2), \dots, X(N)$, the algorithm constructs k new time series:

$$\begin{aligned} X_{km} : X(m), X(m+k), X(m+2k), \dots, \\ X\left(m + \text{int}\left(\frac{N-m}{k}\right) \cdot k\right) \quad \text{for } m = 1, 2, \dots, k, \end{aligned} \quad (12)$$

where m is initial time, k is interval time, $\text{int}(r)$ is integer part of a real number r .

For example, for $k = 4$ and $N = 1000$, the algorithm produces 4 time series:

$$\begin{aligned} X_{41} : X(1), X(5), X(9), \dots, X(997), \\ X_{42} : X(2), X(6), X(10), \dots, X(998), \\ X_{43} : X(3), X(7), X(11), \dots, X(999), \\ X_{44} : X(4), X(8), X(12), \dots, X(1000), \end{aligned} \quad (13)$$

The "length" $L_m(k)$ of each curve X_{km} is then calculated as

$$\begin{aligned} L_m = \frac{1}{k} \left[\left(\sum_{i=1}^{\text{int}((N-m)/k)} |X(m+i \cdot k) - X(m+(i-1) \cdot k)| \right) \right] \\ \times \frac{N-1}{\text{int}\left(\frac{N-m}{k}\right) \cdot k}, \end{aligned} \quad (14)$$

where N is total number of samples.

$L_m(k)$ is not "length" in Euclidean sense, it represents the normalized sum of absolute values of difference in ordinates of pair of points distant k (with initial point m). The "length" of curve for the time interval k , $L(k)$, is calculated as the mean of the k values $L_m(k)$ for $m = 1, 2, \dots, k$:

$$L(k) = \frac{\sum_{m=1}^k L_m(k)}{k}. \quad (15)$$

The value of fractal dimension, D_f , is calculated by a least-squares linear best-fitting procedure as the angular coefficient of the linear regression of the log-log graph of (1):

$$y = ax + b \quad (16)$$

with $a = D_f$, according to the following formulae:

$$D_f = \frac{n \sum (x_k \cdot y_k) - \sum x_k \sum y_k}{n \sum x_k^2 - (\sum x_k)^2}, \quad (17)$$

where $y_k = \ln L(k)$, $x(k) = \ln(1/k)$.

$k = k_1, \dots, k_{\max}$, and n denotes the number of k values for which the linear regression is calculated ($2 \leq n \leq k_{\max}$).

The standard deviation of D_f is calculated as

$$S_{D_f} = \sqrt{\frac{n \cdot \left[\sum y_k^2 - D_f \cdot \sum x_k y_k - b \cdot \sum y_k \right]}{(n-2) \cdot \left[n \cdot \sum x_k^2 - (\sum x_k)^2 \right]}}, \quad (18)$$

where

$$b = \frac{1}{n} \left(\sum y_k - D_f \cdot \sum x_k \right), \quad (19)$$

with standard deviation

$$S_b = \sqrt{\frac{1}{n} \cdot S_{D_f}^2 \cdot \sum x_k^2}. \quad (20)$$

Higuchi's fractal dimension has a scaling feature. Multiplication of all amplitudes X_{km} by a constant factor, c , causes multiplication of the "length" $L_m(k)$ by the same factor. Such multiplication does not change D_f :

$$\ln(L(k)) = D_f \cdot \ln\left(\frac{1}{k}\right) + (b + \ln(c)). \quad (21)$$

Window length has a meaning effect to the results. Because seizures spread so quickly, a displacement as small as possible that does not provide too much variability is desired. We experimented with values ranging from 1 second to 60 seconds and observed that the window length to 2048 points (16 seconds) with 50% overlap should provide reasonable propagation resolution of seizure precursors and the ability of multichannel analysis to effect detection.

3. Results and Discussion

Figure 2 shows an EEG recording of an epilepsy patient which lasted in the vicinity of 21 minutes.

According to the record, it was different between before and after 848th second (14 minutes 08 seconds). Before this point of time, data showed that the neuronal activities were chaotic. However, after that, the brain activity was periodic as a series of high-frequency repetitive spikes. Therefore, it has the ability on seizure onset detection which can probably rely on alteration in fractal characteristic of the signal calculated by Higuchi algorithm. It should be noted that because Higuchi algorithm is so sensitive to noise, preprocessing step should be concentrated on to obtain the most believable results. Therefore, in this study, the preprocessing procedure was carried on by two methods which are described below.

3.1. Method 1. After being analyzed by ICA algorithm, the main component which contains epilepsy wave was illustrated on Figure 3.

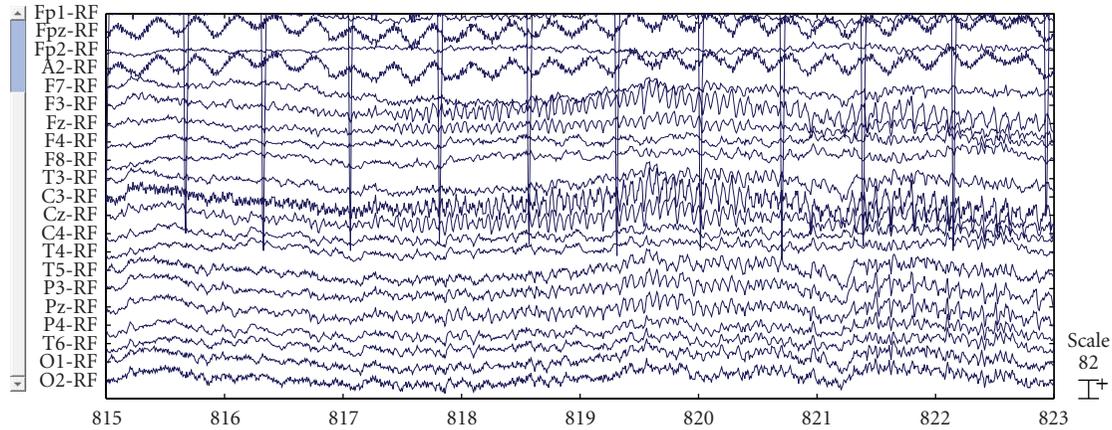


FIGURE 2: The recording of an epilepsy patient.

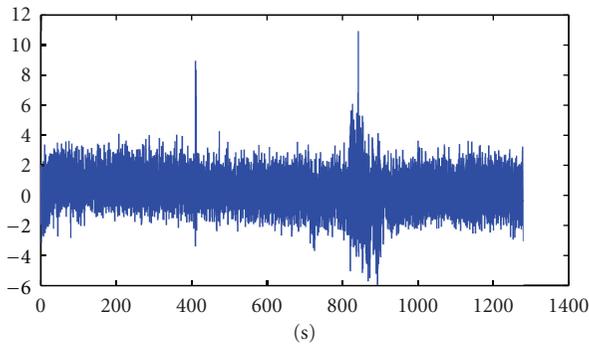


FIGURE 3: The main component containing epilepsy wave following Method 1.

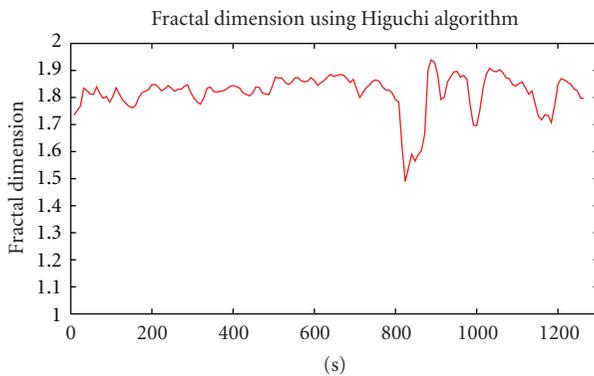


FIGURE 4: Fractal Dimension of IC22 channel.

As can be seen in Figure 3, there were 2 periods of time which had a considerable fluctuation with high amplitude than others. While the first was caused by stimulation effect, the second was the ictal period. The result of FD using Higuchi algorithm is shown in Figure 4.

As regard to Figure 4, the most remarkable aspects of these trends are, during the preictal period, the fractal dimension was relatively high and erratic fluctuates in a small range, hovering at 1.7. This pattern lasted about 13

minutes, until the fractal dimension number reached a peak at 736th second (the window length is 16, overlap 50%). Then the graph declined gradually to 800th second, followed by a sharp fall from 818th second to the trough at 848th seconds. The figure then experienced a recovery, reached to the maximum before falling down to the initial state. The most prominent meaning is that the beginner of ictal period in original data corresponds to the minimal drop in the FD values. Before minimum point occurred approximately in 2 minutes, the fractal dimension value started to decline. Therefore, it is possible to predict some minutes before the happening of seizure.

However, because Higuchi algorithm is very sensitive to noise [2], especially white noise, the average Fractal dimension of each channel in data is so high and it is so difficult to detect epilepsy. The current difficulty is that we cannot know exactly where the main component from results of ICA is processing. Therefore, we propose using averaging filter for the original data. The results of this method will be described clearly below.

3.2. Method 2. According to the Figure 1, the original data experienced two filtering stages before calculating fractal dimension of obtained components. Based on the value of fractal dimension, the results can be separated into two groups of ICA components. Components which had high average fractal dimension value had the same patterns with method 1: during the preictal period the fractal dimension was relatively high and remained stable, the fractal dimension exhibited an substantial decrease during the initial stage of the ictal period, and then it went up again, reaching to a peak, followed by a fall to normal state. Meanwhile, the sign of epilepsy did not appear in the balance group. Therefore, the component which had the highest fractal dimension can be considered as the main components that were showed in Figure 5.

The combination between average filter and ICA brings to us quality results. The main reason is that the advantage of averaging filter can probably reject high frequency components of external noisy source which affected mainly on

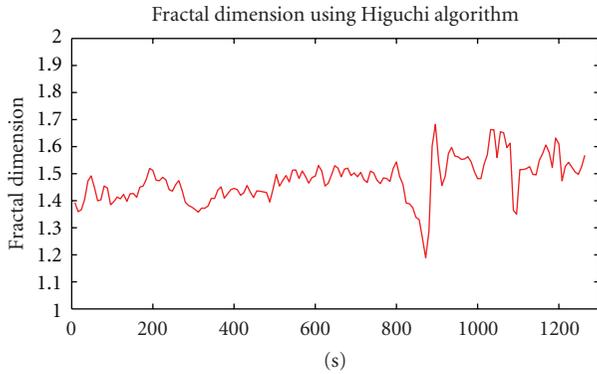


FIGURE 5: The main component containing epilepsy wave following Method 2.

the result of Higuchi algorithm, while ICA is good at rejecting internal noise. From this combination, we can obtain the main component which contained epilepsy waves.

In averaging filter formulae, the length of the window, k , should not be selected too large to lose information of epilepsy wave. This step is suitable for rejecting random noise or very high frequency noise. Therefore, this is an appropriate method in Vietnam condition where equipment, faculty, and measurement condition are not very good. There are not many hospitals applying Faraday cage which is used to eliminate effects of noisy environment.

We noticed that the fractal dimension calculated by Higuchi algorithm has a high degree of accuracy [2]. But, it is very sensitive to noise. So, the step of noise rejection is really important in this research. Using ICA to keep signal separate from noise is not a new way, however, it is so useful in this research. The difficulty when we use this algorithm is that its results include “blind channels”. Therefore, we cannot identify where sources of seizure onset are and which channels have epilepsy wave. The method 2 only helps us to choose which channels to analyze in next step. This is advantage of this method.

The trend of fractal dimension in ictal period has a slight difference from the results of Esteller et al. [7]. Their results showed that the fractal dimension in ictal period is higher than that in the preictal period and ends with a drop to the lowest complexity while the trend of our results obtained an opposite pattern. However, the alteration pattern of the complexity in this study is similar to the results of Iasemidis et al. [11] when they used the Lyapunov exponent for epilepsy data.

3.3. Detect Epilepsy on ECG. Besides achievements in EEG, fractal properties of ECG are also useful for epilepsy diagnosis. While epileptic sign can be visually observed on EEG records, ECG is not paid attention to be considered as a mean playing a substantial role in diagnosing epilepsy. However, in this study, we also attempted to estimate fractal characteristics of ECG of epileptic patients.

According to Figure 6 obtained by the Higuchi algorithm, we can see that fractal coefficient of ECG turned for the worse in the transition from preictal to ictal period. That general

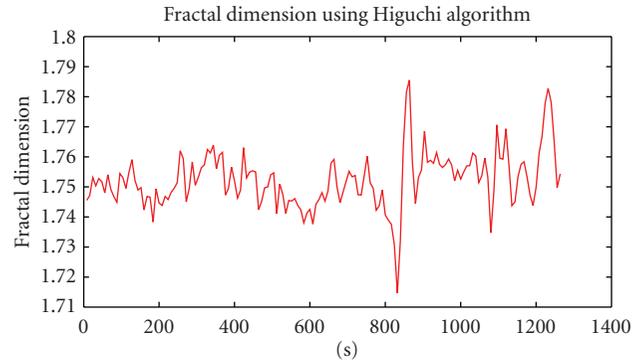


FIGURE 6: The result of fractal dimension on ECG channel.

pattern was very close to the result of EEG when the fall of fractal figure was marked as the beginning of the seizure. In addition, before the seizure by several minutes, there were two troughs that need to be focus on in anticipation of the seizure. That issue had been discussed in study of Iasemidis et al. [11]. That fact makes a proposal that ECG is likely to become a potential method for diagnosis in that domain.

In reality, there is a variety of conveniences of processing ECG in comparison with EEG. Firstly, the former is less sensitive to noise with the great preference for the latter, the main reason is that the amplitude of ECG obtained by the sensors is far higher than that of EEG signal. Secondly, ECG is more widely used than EEG and more suitable for long-term or even perpetual inspection. Therefore, that issue needs to be discussed more deeply because of its advantages.

4. Conclusions

Noise is a serious problem with EEG signal processing, especially in Higuchi algorithm. Therefore, this study concentrated on developing a robust algorithm in the preprocessing step which was the combination between ICA and averaging filter. This fact aimed to reject some kinds of internal and external noise. In addition, this study shows the fractal dimension properties in EEG of epilepsy patients. The results also suggest that FD is a practical tool for identification of seizure onset in the EEG data. The changes in EEG from unperiodic to periodic signal show clearly through the alterations of fractal coefficient to the minimal point. These FD changes may provide insight into the underlying dynamics of this unknown system. These methods can open the possibility of designing an intelligent system for predicting and warning of seizures in real time as a preference or a standard of expert visual analysis of electrographic seizure onset. Moreover, the existing of epileptic sign in fractal result of ECG should be paid attention because of the advantages that could bring to us.

Acknowledgments

The authors are thankful for supports from Department of Science and Technology, Ho Chi Minh City; Vietnam

National Foundation for Science and Technology Development-NAFOSTED Grant No. 106.99-2010.11; Vietnam National University-Ho Chi Minh; They also would like to thank Dr. Cao Phi Phong, Dr. Nguyen Huu Cong, and Dr. Nguyen Thanh Luy for their valuable advices about human anatomy and physiology. Last but not least, they are deeply grateful for the support they have been receiving from their volunteers, families and friends.

References

- [1] N. Pradhan and D. N. Dutt, "Use of running fractal dimension for the analysis of changing patterns in electroencephalograms," *Computers in Biology and Medicine*, vol. 23, no. 5, pp. 381–388, 1993.
- [2] J. Echauz, R. Esteller, B. Litt, and G. Vachtsevanos, "A Comparison of waveform fractal dimension algorithms," *IEEE Transactions on Circuits and Systems I*, vol. 48, no. 2, pp. 177–183, 2001.
- [3] T. Higuchi, "Approach to an irregular time series on the basis of the fractal theory," *Physica D*, vol. 31, no. 2, pp. 277–283, 1988.
- [4] M. J. Katz, "Fractals and the analysis of waveforms," *Computers in Biology and Medicine*, vol. 18, no. 3, pp. 145–156, 1988.
- [5] A. Petrosian, "Kolmogorov complexity of finite sequences and recognition of different preictal EEG patterns," in *Proceedings of the 8th IEEE Symposium on Computer-Based Medical Systems*, pp. 212–217, June 1995.
- [6] N. Kannathal, S. K. Puthusserypady, and L. C. Min, "Complex dynamics of epileptic EEG," in *Proceedings of the 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC '04)*, pp. 604–607, San Francisco, Calif, USA, September 2004.
- [7] R. Esteller, J. Echauz, B. Pless, T. Tchong, and B. Litt, "Real-time simulation of a seizure detection system suitable for an implantable device," *Epilepsia*, vol. 43, supplement 7, p. 46, 2002.
- [8] M. D'Alessandro, R. Esteller, G. Vachtsevanos, A. Hinson, J. Echauz, and B. Litt, "Epileptic seizure prediction using hybrid feature selection over multiple intracranial EEG electrode contacts: a report of four patients," *IEEE Transactions on Bio-Medical Engineering*, vol. 50, no. 5, pp. 603–615, 2003.
- [9] F. S. Bao, D. Y. C. Lie, and Y. Zhang, "A new approach to automated epileptic diagnosis using EEG and probabilistic neural network," in *Proceedings of the 20th IEEE International Conference on Tools with Artificial Intelligence (ICTAI '08)*, pp. 482–486, November 2008.
- [10] M. Ungureanu, C. Bigan, R. Strungaru, and V. Lazarescu, "Independent component analysis applied in biomedical signal processing," *Measurement Science Review. Section 2*, vol. 4, 2004.
- [11] L. D. Iasemidis, J. C. Principe, and J. C. Sackellares, "Measurement and quantification of Spatio-temporal dynamics of human epileptic seizures," in *Nonlinear Biomedical Signal Processing: Dynamic Analysis and Modeling*, vol. 2 of *IEEE Press Series on Biomedical Engineering Metin Akay, Series Editor*, 1999.

Review Article

Linear and Nonlinear Heart Rate Variability Indexes in Clinical Practice

Buccelletti Francesco, Bocci Maria Grazia, Gilardi Emanuele, Fiore Valentina, Calcinaro Sara, Fragnoli Chiara, Maviglia Riccardo, and Franceschi Francesco

Department of Emergency Medicine and Intensive Care, Catholic University of the Sacred Heart, 00168 Rome, Italy

Correspondence should be addressed to Buccelletti Francesco, f.buccelletti@rm.unicatt.it

Received 1 September 2011; Revised 9 November 2011; Accepted 11 November 2011

Academic Editor: Sreenivasan R. Nadar

Copyright © 2012 Buccelletti Francesco et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Biological organisms have intrinsic control systems that act in response to internal and external stimuli maintaining homeostasis. Human heart rate is not regular and varies in time and such variability, also known as heart rate variability (HRV), is not random. HRV depends upon organism's physiologic and/or pathologic state. Physicians are always interested in predicting patient's risk of developing major and life-threatening complications. Understanding biological signals behavior helps to characterize patient's state and might represent a step toward a better care. The main advantage of signals such as HRV indexes is that it can be calculated in real time in noninvasive manner, while all current biomarkers used in clinical practice are discrete and imply blood sample analysis. In this paper HRV linear and nonlinear indexes are reviewed and data from real patients are provided to show how these indexes might be used in clinical practice.

1. Complexity in Biological Signals

Biological systems are complex systems; particularly, they are systems that are spatially and temporally complex, built from a dynamic web of interconnected feedback loops and marked by interdependence, pleiotropy, and redundancy [1]. The meaning of variability in biological signals was studied by Goldberger [2]. He proposed that increased regularity, of signals represents a “decomplexification” of illness. Thus, health is characterized by “organized variability” and disease is defined by decomplexification, increased regularity and reduction in variability. In contrast to the “decomplexification” hypothesis, Vaillancourt and Newell [3] noted increased complexity and increased approximate entropy in several disease states and hypothesized that disease may manifest with increased or decreased complexity, depending on the underlying dimension of the intrinsic dynamic (e.g., oscillating versus fixed point). In addition to the discussion, Macklem's studies on asthma as a disease of higher energy dissipation, greater distance from thermodynamic equilibrium, lower entropy, and greater variation [4] suggest that health is defined by a certain distance from thermodynamic

equilibrium; too close (decreased variation, too little energy dissipation, low entropy) or too far (increased variation and energy dissipation, high entropy) each represents pathological alterations [5].

The host response to sepsis, shock, or trauma is an example of a biological complex system that is readily apparent to intensivists [6]. It is within this complex systems conception of health and illness that the clinical utility of variability analysis may be appreciated and should determine the impact that the variability analysis has on critically ill patient outcome.

If we look at a modern emergency department (ED) and intensive care unit (ICU) we can appreciate a continuous stream of information: parameters derived by multiple monitors and ventilators, laboratory data, and clinical documentation. Usually, data are collecting intermittently but this system is not adequate for tracking and analysis of complex multivariate relationships. Variability analysis represents a novel means to evaluate and treat individual patients, suggesting a shift from epidemiological analytical investigation to continuous individualized complexity analysis [7]. Complexity analysis of time series has been widely

used in the study of variability of biological phenomena, as heart rate [8].

Heart rate is probably the easiest biological, complex, signal to analyze. Heart rate, recorded as a space between two heartbeats or as a distance R-R on an surface electrocardiogram (ECG), is irregular if measured in milliseconds. This kind of variation appeared significant and is related to physiological (or pathological) conditions. Previous studies demonstrated a fractal-like complexity pattern in the variability of heart rate (HRV) which is possible to measure and quantify. Rapid fluctuation of HRV can reflect changes of sympathetic and parasympathetic activity; in other words, HRV is a noninvasive index of the autonomic nervous system's control on the heart. Recent studies suggested that mechanisms involved in the regulation of cardiovascular system interact with each other in a nonlinear way and that it is possible to study these mechanisms with several algorithms. Clinically, patients after an acute myocardial infarction showed altered HRV indexes values with such differences correlating to overall mortality [9].

The aim of this paper is to describe different approaches to HRV quantifications in real patients, all of possible utility in future clinical practice.

2. Heart Rate Variability Indexes

See references for a clear and exhaustive explanation of different HRV indexes and their meanings and clinical use [10]. Table 2 shows most used HRV indexes in clinical practice.

2.1. Linear Algorithms. Using linear algorithms, HRV can be analyzed in time or frequency domain. *Time domain indexes* are the first used indexes and simplest way to calculate HRV, because they are statistical calculations of consecutive RR intervals, and they are strictly correlated with each other (SDNN, SDANN, pNN50, ecc...). *Frequency domain indexes* are more elaborated indexes based on spectral analysis, mostly used to evaluate the contribution on HRV of autonomic nervous system (VLF, LF, HF, HF/LF ratio). Spectral analysis can be used to analyze the sequence of NN intervals of short-term recordings (2 to 5 minutes) or an entire 24-hour period (i.e., Holter-ECG record).

2.2. Nonlinear Algorithm. Non linear (fractal) indexes [11] are recently introduced methods to measure HRV, not affected by nonstationarity, as it happens for linear indexes. They include Power Law Exponent, Approximate Entropy and Detrended Fluctuation Analysis. These methods study all complex interactions of hemodynamic, electrophysiological, and humoral variables as well as by the autonomic and central nervous regulations. These techniques have been shown to be powerful tools for characterization of various complex systems, but however no systematic study has been conducted to investigate large patient populations with the use of these methods.

Starting from frequency analysis, *Power Law Exponent* [12] describes the nature of correlations of single frequencies in a time series. When equal to 1, it states that the time

series has similar fluctuations acting at different scales, regardless of the size of the variation (namely, it is "scale invariant," a property of fractals [13]). It has been applied in biology and medicine formerly to describe the dynamics of beat-to-beat interval in ageing [8]. *Approximate Entropy* (ApEn) [14] provides a measure of the degree of irregularity or randomness within a series of data. Smaller values indicate greater regularity, and greater values convey more randomness and system complexity. It is a rather new index applied in biological systems signals study and it still needs implementation.

2.3. Detrended Fluctuation Analysis (DFA). This method has been developed in order to make a distinction between the internal variations generated by complex systems and those variations caused by some environmental-external stimulus [15]. A singular ECG derivation is recorded continuously and the R-R distance is calculated in milliseconds until it is possible to get an amount of 8000 R-R that are necessary to assure an adequate interval of time. The data's series have been integrated and divided into a series of regular intervals named n , included between 1 and 300. For each n interval, it has been calculated the "local" fluctuation as the difference compared to a straight line of a linear interpolation. Indeed, the "global" fluctuation has been calculated as the square root of the average of the local's fluctuations.

3. Examples of HRV Indexes in Healthy and Critically Ill Patients

HRV indexes were computed using a digital 12 leads ECG-Holter machine (Mortara Instruments, USA) in twenty consecutive patients admitted to the Intensive Care Unit within 24 hours of admission and in an aged-matched (2:1 ratio) control population from consecutive patients presenting to the ED with nontraumatic, self-limited, chest pain as chief complain. They were at very low/low/medium risk [16] for Acute Coronary Syndrome (ACS), and entered in a dedicated protocol to be screened for silent cardiac ischemia with serial cardiac enzymes measures (high-sensitivity Troponin T, Elecsys, Roche, Germany) followed by provocative cardiac stress test (either treadmill or nuclear stress test). All patients had to have at least 18 years of age and a baseline 12 leads ECG without diagnostic T-wave or ST-segment deviation suggesting ongoing acute coronary syndrome. Study protocol was approved by Ethical Committee. ECGs were manually reviewed and only patients with sinus rhythm at baseline were eligible for HRV indexes computation. ECG-Holter data were analyzed by Mortara proprietary software to obtain RR intervals in milliseconds, then HRV indexes were manually calculated. Holter data with artifacts or nonsinus beat more than 10% of total beats number were excluded from the analysis.

Critically ill patients consisted in a mixed population treated in the Medical Intensive Care Unit (MICU) and in the Surgical Intensive Care Unit (SICU). The SICU group enrolled patients after major heart surgery (coronary artery bypass graft or valvular replacement) while the MICU

TABLE 1: Descriptives. All variables are displayed as mean (standard deviation) except gender expressed as number (%).

	Critically Ill Patients ($n = 20$)	ED Patients ($n = 45$)
Age (years)	54 (20)	57 (14)
Gender (Male)	15 (75)	25 (55)
SAPS	38 (17)	—
SOFA	6.3 (4.6)	—
APACHE II	15 (7)	—
SBP (mmHg)	132 (23)	140 (23)
DBP (mmHg)	73 (24)	80 (20)
Heart Rate (bpm)	81 (19)	80 (19)
SDNN index (ms)	40 (21)	53 (34)
LF/HF Ratio 5 min (ms^2)	0.97 (0.67)	0.84 (0.34)
LF/HF Ratio tot. (ms^2)	0.89 (0.40)	0.85 (0.23)
DFA		
Alpha1	0.76 (0.43)	0.98 (0.31)
Alpha2	0.99 (0.18)	1.01 (0.09)

ED: emergency Department. SBP: systolic blood pressure. DBP: diastolic blood pressure. SDNN: standard deviation of the NN intervals. LF: low frequency. HF: high frequency. DFA: detrended fluctuation analysis.

patients were treated for septic shock (defined as infection in the setting of high serum lactate and unstable hemodynamic conditions at presentation). All patients were mechanically ventilated and under treatment with Propofol.

In a study published in 2005 on Anesthesia Analgesia, Propofol induces significant decreases in BP, LF, HF, ApEn, and LF/HF ratio with no change in HR, indicating predominance of parasympathetic activity during sedation. The decreased BP with no change in HR indicates that propofol attenuates the baroreflex reaction [17]. Kanaya et al. [18] reported that continuous infusion of propofol at a rate of 3 mg/kg/min reduced cardiac parasympathetic tone based on a decrease in entropy and HF with no significant changes in LF, LF/HF, and HR.

All patients were lying supine and no invasive procedures were performed during ECG signal acquisition. ECG signal was sampled at 1 KHz (1000 samples/sec), assuring a good quality of measurement in a millisecond scale. STATA 11.0 (Stata Corp, TX) was used to compute statistics. Table 1 provides clinical characteristics of the two groups. Figures 1, 2, and 3 show the three most used HRV indexes behaviours in the two groups.

4. Bridging the Gap: From Research to Clinical Practice

As clinicians, ED physicians and intensivists are always interested in predicting patient's risk of developing major and life-threatening complications. Such risk models should include biomarkers of prognostic value. Anticipating clinical course implies a deep knowledge of the present patient's state.

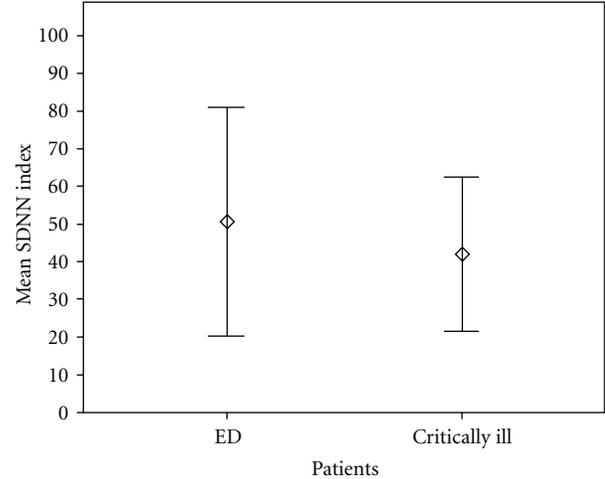


FIGURE 1: Series Standard Deviation (Frequentist Statistics). SDNN index displayed as mean (circles) and 95% confidential interval (Bars). Healthy subjects showed a higher degree of dispersion around the mean (higher variability) compared to critically ill patients, $P = 0.10$ using Mann-Witney U -test. ED: Emergency Department.

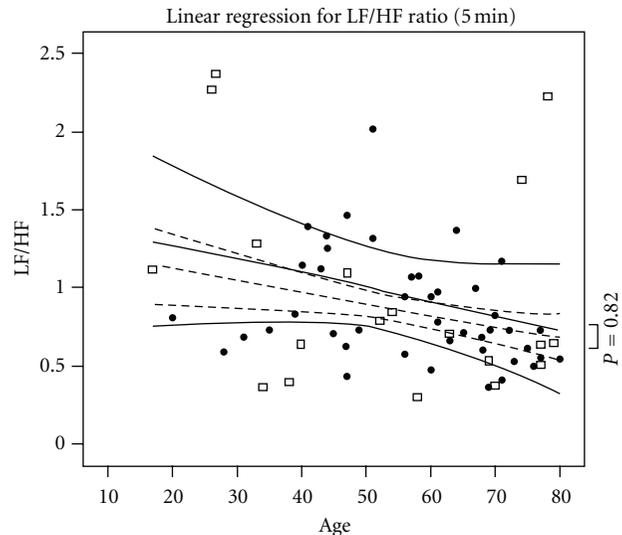


FIGURE 2: Fast Fourier Transform Analysis. Black dots represent healthy patients and empty squares ICU cases. Dashed and continuous lines reflect LF/HF ratio after adjusting for other clinical comorbidities along with 95% confidential intervals (curved lines). x -axis represents age in years. The two groups did not differ in term of LF/HF ratio ($P = 0.82$).

Understanding biological signals behavior might represent a fundamental step toward a better care. The main advantage of signals such HRV indexes is that it can be calculated in real time in noninvasive manner. In fact, all current biomarkers used in clinical practice are discrete and imply blood sample analysis.

Although HRV indexes appear to be appealing, further research is required. First, nonlinear indexes are not standardized in terms of data gathering methodology and

TABLE 2: Most used HRV measures. (modified from *Heart Rate Variability: Standards of Measurement, Physiological Interpretation, and Clinical Use. Task Force of the European Society of Cardiology the North American Society of Pacing Electrophysiology. Circulation. 1996;93:1043–1065*).

Time domain indexes	SDNN	Standard deviation of all NN intervals
	SDANN	Standard deviation of the averages of NN intervals in all 5-minute segments of the entire recording
	RMSSD	The square root of the mean of the sum of the squares of differences between adjacent NN intervals
	SDNN index	Mean of the standard deviations of all NN intervals for all 5-minute segments of the entire recording
	pNN50	NN50 count divided by the total number of all NN intervals
Frequency domain indexes	Total power	Variance of all NN intervals (≤ 0.4 Hz)
	ULF	Power in the ULF range (≤ 0.003 Hz)
	VLF	Power in the VLF range (0.003–0.04 Hz)
	LF	Power in the LF range (0.04–0.15 Hz)
	HF	Power in the HF range (0.15–0.4 Hz)
	LF/HF	Ratio LF [ms^2]/HF [ms^2]

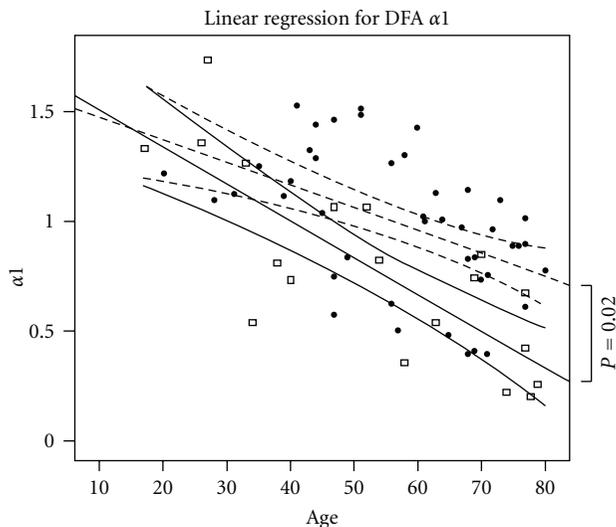


FIGURE 3: Detrended Fluctuation Analysis (DFA). Black dots represent healthy patients and empty squares ICU cases. Dashed and continuous lines reflect predicted values (adjusted for comorbidities) for the respective group along with 95% confidential intervals (curved lines). x -axis represents years. It is to be noted that DFA index was significantly different between the two groups even when adjusted for other comorbidities and age ($P = 0.02$). Age affects DFA index in both groups.

minimum numbers of R-R intervals needed to have back a reliable measure. Second, HRV indexes represent the final outcome of complex systems. For instance, it is known that diabetes significantly affects final results, along with aging. All those chronic clinical characteristics are not well studied and no nomograms exist to simply adjust indexes results for these covariates. Third, it is not clear how and why different algorithms behave in different manner. For instance, fast Fourier analysis was shown to provide information and was able to discriminate between patients with and without coronary artery disease, but it seemed not to be the same case in our example (although sample size is similar in our case and in previous study [19]).

5. Conclusions

Different pathophysiologic processes alter HRV indexes in opposite directions, making it difficult to identify them when present at the same time. For instance, age and diabetes both decrease DFA index while acute myocardial ischemia seems to increase it. HRV-indexes future studies should be aimed to evaluate how HRV is affected by known cardiovascular risk factors and to find a “standard” of measurement of different indexes, comparing healthy and ill patients and investigating their risk of major cardiovascular events.

In conclusion, future larger studies are warranted before HRV indexes can be embedded into daily clinical practice as routine standard of care.

References

- [1] R. Gallagher and T. Appenzeller, “Beyond reductionism,” *Science*, vol. 284, no. 5411, p. 79, 1999.
- [2] A. L. Goldberger, “Fractal variability versus pathologic periodicity: complexity loss and stereotypy in disease,” *Perspectives in Biology and Medicine*, vol. 40, no. 4, pp. 543–561, 1997.
- [3] D. E. Vaillancourt and K. M. Newell, “Changing complexity in human behavior and physiology through aging and disease,” *Neurobiology of Aging*, vol. 23, no. 1, pp. 1–11, 2002.
- [4] C. L. Que, C. M. Kenyon, R. Olivenstein, P. T. Macklem, and G. N. Maksym, “Homeokinesis and short-term variability of human airway caliber,” *Journal of Applied Physiology*, vol. 91, no. 3, pp. 1131–1141, 2001.
- [5] A. J. Seely and P. T. Macklem, “Complex systems and the technology of variability analysis,” *Critical Care*, vol. 8, no. 6, pp. R367–R384, 2004.
- [6] A. J. E. Seely and N. V. Christou, “Multiple organ dysfunction syndrome: exploring the paradigm of complex nonlinear systems,” *Critical Care Medicine*, vol. 28, no. 7, pp. 2193–2200, 2000.
- [7] M. J. Cohen, A. D. Grossman, D. Morabito, M. M. Knudson, A. J. Butte, and G. T. Manley, “Identification of complex metabolic states in critically injured patients using bioinformatic cluster analysis,” *Critical Care*, vol. 14, no. 1, article R10, 2010.
- [8] S. M. Pikkujämsä, T. H. Mäkikallio, L. B. Sourander et al., “Cardiac interbeat interval dynamics from childhood to

- senescence: comparison of conventional and new measures based on fractals and chaos theory,” *Circulation*, vol. 100, no. 4, pp. 393–399, 1999.
- [9] F. Buccelletti, E. Gilardi, E. Scaini et al., “Heart rate variability and myocardial infarction: systematic literature review and meta-analysis,” *European Review for Medical and Pharmacological Sciences*, vol. 13, no. 4, pp. 299–307, 2009.
- [10] M. Malik, “Heart rate variability: standards of measurement, physiological interpretation, and clinical use,” *Circulation*, vol. 93, no. 5, pp. 1043–1065, 1996.
- [11] J. M. Tapanainen, P. E. B. Thomsen, L. Køber et al., “Fractal analysis of heart rate variability and mortality after an acute myocardial infarction,” *American Journal of Cardiology*, vol. 90, no. 4, pp. 347–352, 2002.
- [12] T. Gisiger, “Scale invariance in biology: coincidence or footprint of a universal mechanism?” *Biological Reviews of the Cambridge Philosophical Society*, vol. 76, no. 2, pp. 161–209, 2001.
- [13] B. Mandelbrot, *The Fractal Geometry of Nature*, W. H. Freeman, New York, NY, USA, 1983.
- [14] S. M. Pincus, “Approximate entropy as a measure of system complexity,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 88, no. 6, pp. 2297–2301, 1991.
- [15] C. K. Peng, S. Havlin, H. E. Stanley, and A. L. Goldberger, “Quantification of scaling exponents and crossover phenomena in nonstationary heartbeat time series,” *Chaos*, vol. 5, no. 1, pp. 82–87, 1995.
- [16] H. V. Huikuri and T. H. Mäkikallio, “Heart rate variability in ischemic heart disease,” *Autonomic Neuroscience*, vol. 90, no. 1-2, pp. 95–101, 2001.
- [17] N. N. Win, H. Fukayama, H. Kohase, and M. Umino, “The different effects of intravenous propofol and midazolam sedation on hemodynamic and heart rate variability,” *Anesthesia and Analgesia*, vol. 101, no. 1, pp. 97–102, 2005.
- [18] N. Kanaya, N. Hirata, S. Kurosawa, M. Nakayama, and A. Namiki, “Differential effects of propofol and sevoflurane on heart rate variability,” *Anesthesiology*, vol. 98, no. 1, pp. 34–40, 2003.
- [19] K. Sroka, C. J. Peimann, and H. Seevers, “Heart rate variability in myocardial ischemia during daily life,” *Journal of Electrocardiology*, vol. 30, no. 1, pp. 45–56, 1997.

Research Article

A Finite Element Study of Micropipette Aspiration of Single Cells: Effect of Compressibility

Amirhossein Jafari Bidhendi and Rami K. Korhonen

Department of Applied Physics, University of Eastern Finland, P.O. Box 1627, 70211 Kuopio, Finland

Correspondence should be addressed to Rami K. Korhonen, rami.korhonen@uef.fi

Received 27 September 2011; Accepted 31 October 2011

Academic Editor: Vikas Rai

Copyright © 2012 A. Jafari Bidhendi and R. K. Korhonen. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Micropipette aspiration (MA) technique has been widely used to measure the viscoelastic properties of different cell types. Cells experience nonlinear large deformations during the aspiration procedure. Neo-Hookean viscohyperelastic (NHVH) incompressible and compressible models were used to simulate the creep behavior of cells in MA, particularly accounting for the effect of compressibility, bulk relaxation, and hardening phenomena under large strain. In order to find optimal material parameters, the models were fitted to the experimental data available for mesenchymal stem cells. Finally, through Neo-Hookean porohyperelastic (NHPH) material model for the cell, the influence of fluid flow on the aspiration length of the cell was studied. Based on the results, we suggest that the compressibility and bulk relaxation/fluid flow play a significant role in the deformation behavior of single cells and should be taken into account in the analysis of the mechanics of cells.

1. Introduction

Recent studies have shown that the mechanical factors play an important role in cell metabolism, differentiation, and function [1, 2]. Cells possess mechanical properties that change along with changes in cytoskeletal network in each stage of differentiation or disease [3]. Mechanobiology of tissues and cell biosynthesis are related to the mechanical signals experienced by cells, and the deformation behavior of cells within a tissue is affected by the mechanical environment and mechanical properties of cells. Hence, a method to search for the appropriate cells in a mass of cultured stem cells may be based on measuring the mechanical properties of those cells. Moreover, evaluation of the mechanical properties of cells improves the understanding of the onset and progression of tissue pathologies.

Micropipette aspiration (MA) technique has been widely used to measure the mechanical properties of different cell types [1, 4–9]. In this experiment, a cell is aspirated into a micropipette by exerting a negative pressure gradient, and the aspiration length of the cell inside the micropipette is recorded as a function of time. In MA of many cell types, the

cell exhibits an initial jump inside the micropipette under a constant pressure, which is then followed by a creep until reaching the equilibrium. The time-dependent deformation behavior of the cell cannot be modeled with an elastic model. An analytical solution presented by Sato et al. [4] has been commonly employed to interpret the viscoelastic parameters of the cells from experimental data considering a three-parameter solid for the cell (Figure 1(a)). In this model, the cell is assumed as an infinite, homogenous, incompressible half-space under a uniform pressure gradient (Figure 1(b)). Small strain tensors are used for the cell deformation inside the micropipette. Although this solution has been normally used to interpret the aspiration data for its simplicity, there are some limitations in this model. The semi-infinite space assumption for the cell is only accurate when the cell to micropipette radii is considerably large, and the cell is assumed as an incompressible material, the validity of which should be scrutinized. Additionally, in MA cells undergo very large strains and cells exhibit a strain hardening behavior [10, 11], which are not included in this analytical solution.

Hyperelastic models are able to account for nonlinear behavior of cells under large strains. The time-dependent

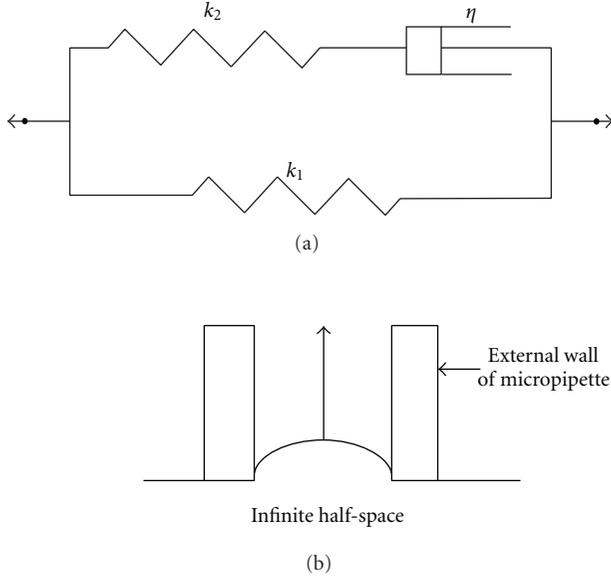


FIGURE 1: (a) Three-parameter viscoelastic linear solid for the cell and (b) infinite half-space model for the cell in MA.

deformation behavior of different cell types has been modeled using viscoelastic, poroelastic and poroviscoelastic material models [8, 12–22]. Zhou et al. [9] studied the influence of micropipette and cell geometries on the aspiration length of the cell inside the micropipette using a Neo-Hookean viscohyperelastic (NHVH) material. However, this model neglects the importance of cell compressibility, which results in different mechanical parameters estimated for the cell. Furthermore, different assumptions of compressibility and viscoelastic properties for the cells will lead to different patterns and magnitudes of stress and strain fields within the cell and cytosol when studying the coupling of external forces in cell-extracellular matrix (ECM) interaction. The compressibility may also be an indicator of the level of cell's cytoskeletal network integrity [23]. In various other studies, the cell has been either considered as a fully incompressible material or has been assumed compressible without any emphasis on its importance [8]. Furthermore, the equivalence or the similarity of the compressibility and viscoelastic bulk relaxation and the fluid flow in the nonlinear deformation behavior of the cells should be studied.

In the current study, we developed a finite element (FE) model of the cell aspiration by applying the compressible NHVH material model. Material parameters of the model were optimized by fitting the model to the experimental data of the MA of mesenchymal stem cells [5]. We investigated the effect of different material parameters and especially the cell compressibility on the creep response of the cells in MA. Finally, Neo-Hookean porohyperelastic (NHPH) material model was developed for the cell to investigate the influence of fluid flow in MA and to clarify if the compressibility and bulk relaxation in the NHVH model could be explained by intracellular fluid.

2. Materials and Methods

2.1. Material Models. The NHVH solid was used to particularly account for the nonlinear mechanical response of cells at large strains and subsequent strain-hardening phenomenon. The Neo-Hookean material model uses a general strain energy potential [24, 25],

$$U(\bar{I}_1, J) = C_{10}(\bar{I}_1 - 3) + \frac{1}{D_1}(J - 1)^2, \quad (1)$$

where U is the strain energy per unit of reference volume, C_{10} and D_1 are temperature-dependent material parameters, and \bar{I}_1 is the first invariant of the deviatoric left Cauchy-Green deformation tensor \bar{B} , defined as $\bar{B} = \bar{F} \cdot \bar{F}^T$. Here \bar{F} is the distortion gradient, $\bar{F} = J^{1/3} F$, and J is the elastic volume ratio. The constitutive equation for the Neo-Hookean material is given by

$$\sigma_s = \frac{2}{J} C_{10} \left(\bar{B} - \frac{1}{3} \text{tr}(\bar{B}) I \right) + \frac{2}{D_1} (J - 1) I, \quad (2)$$

where σ_s is the Cauchy stress. The material parameters are given by

$$C_{10} = \frac{G_0}{2}, \quad D_1 = \frac{3(1 - 2\nu)}{G_0(1 + \nu)}, \quad (3)$$

where G_0 is the initial shear modulus and ν is Poisson's ratio.

Viscoelastic properties can be defined with a Prony expansion of the dimensionless shear and bulk relaxation moduli. The Prony expansion of the time-dependent shear behavior can be formulated as [24]

$$g_R(t) = 1 - \sum_{i=1}^N \bar{g}_i^p \left(1 - e^{-t/\tau_i^G} \right), \quad (4)$$

where $g_R(t)$ ($G_R(t)/G_0$) is the dimensionless shear relaxation modulus, $G_R(t)$ is the shear modulus at the time t , N is the number of terms in the Prony series, \bar{g}_i^p is the dimensionless Prony series parameter for shear modulus, and τ_i^G is the relaxation characteristic time. Several studies have included the shear relaxation behavior to express the time-dependent changes in cell stiffness [4–6, 9]. Similarly, the time-dependent bulk behavior of the material can be expressed in the form of

$$k_R(t) = 1 - \sum_{i=1}^N \bar{k}_i^p \left(1 - e^{-t/\tau_i^K} \right), \quad (5)$$

where $k_R(t)$ is the dimensionless bulk relaxation modulus, τ_i^K is the relaxation characteristic time, and \bar{k}_i^p is the dimensionless Prony series parameter for bulk modulus. The bulk relaxation behavior reflects the time-dependent changes in cell volumetric behavior, and it has also been recently included in modeling the cell compression [16].

Finally, the NHPH model with a Neo-Hookean hyperelastic solid matrix, defined by C_{10s} and D_{1s} , fully saturated with the intracellular fluid was developed. According to the biphasic theory, the total stress in the cell is defined as [26]

$$\sigma_t = -pI + \sigma_s, \quad (6)$$

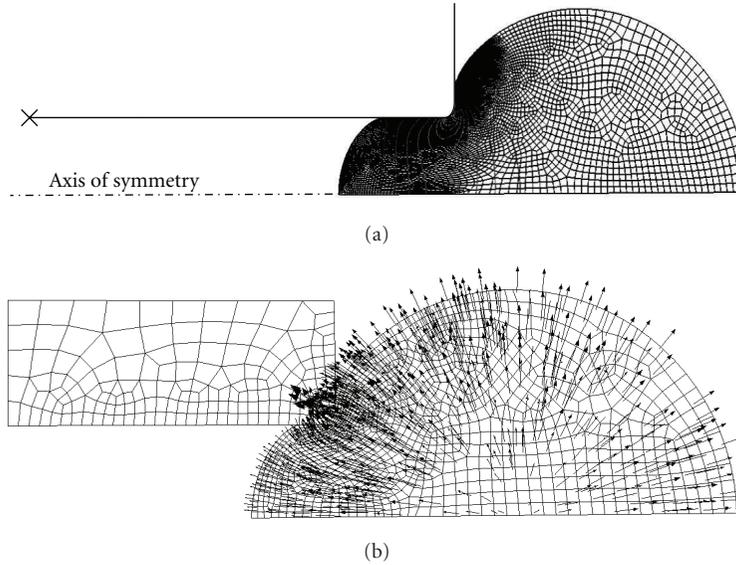


FIGURE 2: (a) The axisymmetric FE mesh of micropipette aspiration. (b) NHPH model of micropipette aspiration simulating the fluid exchange and velocity in boundary of the cell.

where p is the hydrostatic pressure, σ_s is the effective solid stress tensor which is derived from (2) and I is the unit tensor. Fluid flow in the cell was modeled according to Darcy's law as [24]

$$q = -k\nabla P, \quad (7)$$

where q is the flow rate, k is the permeability (with the unit m^4/Ns), and ∇P is the pressure gradient across the region.

2.2. Finite Element Analysis. In the NHVH FE model of the MA, axisymmetric geometries were created for the micropipette and cell, resembling the 3-dimensional, spherical shape of the cell in suspension. The cell was considered as a homogenous continuum discretized with four-node bilinear quadratic hybrid axisymmetric elements, CAX4H (Figure 2(a)). In critical contact areas, the number of elements was increased. The micropipette was considered as analytical rigid, since it necessitates less computation time. A fillet radius was considered at the opening of the micropipette to mimic the experimentally used micropipettes [5, 9]. For a sufficiently large micropipette, such as the one used in the present study, the effect of fillet radius on the modeled response of the cell has been shown to be insignificant [9, 27]. Symmetry boundary condition was employed on the cell, disallowing the horizontal movement of the cell in the axis of symmetry. The micropipette was fixed at its reference point. Frictionless surface-to-surface contact (Abaqus/Standard, Dassault Systèmes Simulia Corp, Providence, RI, USA) was implemented between the micropipette and the cell surface. Since the inertial forces are negligible during the aspiration, the procedure was considered quasistatic. The aspiration pressure was set to 890 Pa, according to the experiments, reaching its maximum almost instantaneously

(within 0.001 s). Abaqus 6.8.1 finite element package was used in all pre- and postprocessing. Sensitivity of the results to the mesh density and the element type was investigated.

For the NHPH model, in addition to the previous boundary conditions, the pore pressure at the external surface of the cell was set to zero to allow free fluid flow through the boundary. In this model, general contact formulation (Abaqus/Standard) was envisaged between the micropipette and the cell surface to allow free fluid flow in the sliding interaction boundary. Thus, the micropipette was modeled as a linearly elastic solid with a very high Young's modulus to resemble the rigid glass (in order to ensure the reliability of the results in both of the contact formulations, the FE models with elastic and rigid micropipettes were compared by means of the NHVH material for the cell. The results showed a negligible difference). Four-node axisymmetric pore fluid/stress finite elements (CAX4P) were used (Figure 2(b)), and the simulations were conducted by applying the soils consolidation analysis.

2.3. Optimization. To find optimum set of material parameters for the models, they were fitted to the experimental creep data [5]. For that, direct search or particle swarm optimization (PSO) methods were utilized depending on the number of optimized parameters. In the case of an incompressible NHVH model, a direct search method was employed to find the optimum material parameters. In compressible NHVH models, a PSO algorithm was used based on a code between MATLAB (R2008a, MathWorks Inc., Natick, MA, USA) and Abaqus for automatic iterative minimization of weighted sum of squared errors between experimental data points and their corresponding FE model outputs. The efficiency of this algorithm in an FE application was previously confirmed [28]. In order to reduce the optimization time, the ranges for the optimized parameters

TABLE 1: Ranges of the material parameters allowed in optimization.

Parameter	Range
Characteristic time (τ)	1–7 (s)
Shear relaxation parameter (\bar{g}^p)	0.35–0.52 (—)
Bulk relaxation parameter (\bar{k}^p)	0.35–0.73 (—)
Initial shear modulus (G_0)	319–420 (Pa)
Poisson’s ratio (ν)	0.3–0.5 (—)

TABLE 2: Ranges of the material parameters in the parametric study.

Parameter	Range
Shear relaxation parameter (\bar{g}^p)	0.2–0.99 (—)
Bulk relaxation parameter (\bar{k}^p)	0–0.8 (—)
Poisson’s ratio (ν)	0.3–0.5 (—)
Permeability (k)	1×10^{-20} – 1×10^{-11} (m^4/Ns)
Poisson’s ratio of poroelastic matrix (ν_s)	0.01, 0.2, 0.35, 0.42 (—)

were limited (based on the estimated values for NHVH and analytical models, Table 1).

2.4. Parametric Study. Parametric studies were carried out to demonstrate the influence of different material parameters on the response of the cell in MA. Specifically, the effect of compressibility on the deformation behavior of the cell was simulated by varying Poisson’s ratio from 0.3 to 0.5, which is a commonly reported range in the literature for different cell types [7, 8, 29, 30].

In the NHPH model, a range of permeabilities ($k = 1 \times 10^{-20} - 1 \times 10^{-11}$ (m^4/Ns)) were chosen for the analysis (Table 2), covering the minimum and maximum values reported for the cell [31] and membrane [32, 33]. In the study of the contribution of fluid flow to the creep deformation of the cell, the aspiration pressure was reduced to 445 Pa to decrease the computation times. Further, the effect of Poisson’s ratios (0.01, 0.2, 0.35, and 0.42, Table 2), in addition to the range of permeabilities, on creep were tested. The shear modulus was constant in these simulations (319 Pa, optimized initial shear modulus for incompressible NHVH model, Table 3).

3. Results

3.1. Optimization. The compressible NHVH model captured well the experimental creep behavior of cells, while the incompressible NHVH model underestimated the experimental curve, especially the initial time points (Figure 3, Table 3). The incompressible NHVH model with the material parameters obtained from the literature [5] based on the Sato et al. [4] analytical solution underestimated the creep response of cells even more (Figure 3, Table 3). As compared to the incompressible NHVH model, the analytical solution overestimated the characteristic time by nearly 147%. It also underestimated the initial and infinite shear moduli by 7.5%

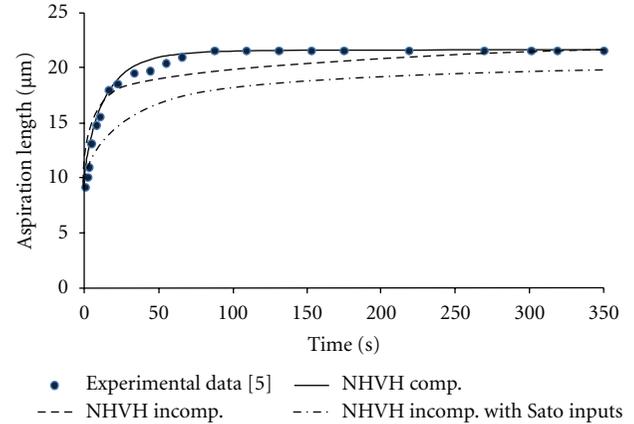


FIGURE 3: Optimized compressible and incompressible NHVH models as well as incompressible NHVH with model inputs from the literature [5] based on the Sato et al. analytical solution.

and 1.5%, respectively (Table 3). Characteristic time, initial and infinite shear moduli differentiated by 100%, 26%, and 47% between the incompressible and compressible NHVH models, respectively (Table 3).

3.2. Parametric Study. In the NHVH model, the initial and final aspiration lengths increased by 48% and 49%, respectively, by the change of Poisson’s ratio from 0.5 to 0.3, while the time to reach the equilibrium remained constant ($K(0) = K(\infty)$, Figure 4). By increasing \bar{g}^p (4), the cell behavior approached a Maxwell droplet ($g_R(\infty) = 0$); the aspiration length was increased and it took longer for the cell to reach the equilibrium length (Figure 5). By increasing \bar{k}^p (5), the length of aspiration was increased, while the equilibrium time remained unchanged (Figure 6). With lower Poisson’s ratios, the bulk relaxation had more influence on the creep deformation of the cell (zero at $\nu = 0.5$) (Figure 7). Furthermore, with lower Poisson’s ratios, the required shear modulus to reproduce the same initial elastic jump became higher, and in contrast, a lower corresponding \bar{k}^p was needed to follow the same bulk relaxation creep.

The effect of fluid flow on the creep behavior of cells in MA was amplified by the decrease in the permeability and Poisson’s ratio, as was shown by the NHPH model (Table 4). However, the creep deformation reached a plateau and was virtually the same with the permeability values lower than 1×10^{-15} (m^4/Ns).

4. Discussion

In this study, creep behavior of a single cell in micropipette aspiration was modeled by applying incompressible and compressible NHVH material models. The model including compressibility (Poisson’s ratio < 0.5) was able to capture the entire experimental aspiration curve. Although the incompressible NHVH model was able to mimic the long-term viscoelastic behavior of the cell, it failed to capture the early creep data points. Consequently, it was suggested that

TABLE 3: Optimized material parameters of the models. The values of root mean square error (RMSE) between experimental data points and corresponding points of each model are also listed.

Model	ν	τ (s)	\bar{k}^P	\bar{g}^P	G_0 (Pa)	RMSE
Incompressible NHVH with inputs from Tan et al. [5]	0.5	7.8	0	0.58	296	2.29
Incompressible NHVH	0.5	1.2	0	0.61	319	1.42
Compressible NHVH	0.42	3.6	0.7	0.51	414	0.44

Parameters for incompressible and compressible NHVH models were obtained by optimizing the FE model to the corresponding experimental curve, while the parameters for the incompressible NHVH model were also obtained from the literature [5].

TABLE 4: Influence of poroelasticity in creep of the cell into micropipette in different Poisson's ratios ($G_0 = 319$ Pa and aspiration pressure = 445 Pa).

Poisson's ratio	Equilibrium aspiration length (μm)	% Poro. at $k = 1 \times 10^{-11}$ (m^4/Ns)	% Poro. at $k = 1 \times 10^{-15}$ (m^4/Ns)	% Poro. at $k = 1 \times 10^{-20}$ (m^4/Ns)
0.42	6.1	3.2	16.8	17.0
0.35	7.0	8.2	27.4	27.6
0.20	8.9	20.0	43.2	43.2
0.01	11.6	33.2	56.0	56.1

% Poro. = ((equilibrium aspiration length – initial aspiration length)/equilibrium aspiration length) \times 100.

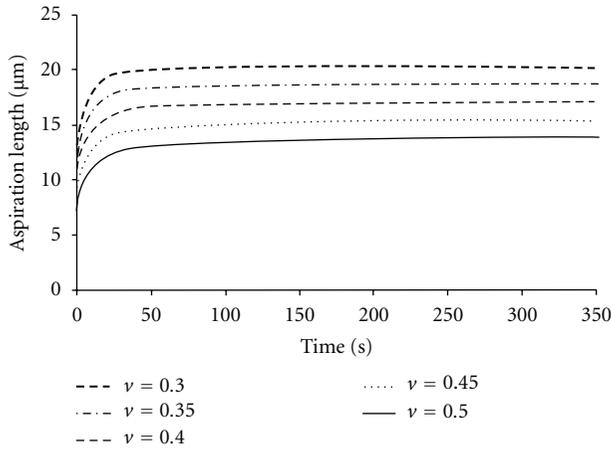


FIGURE 4: The FE results for NHVH cell model with different values of Poisson's ratio (ν) without bulk relaxation ($K_\infty = K_0$). Other parameters were Prony shear relaxation parameter $\bar{g}^P = 0.51$, time constant $\tau = 3.6$ (s), initial shear modulus $G_0 = 414$ (Pa), and aspiration pressure $P = 890$ (Pa).

the creep behavior of the cell can be attributed to both shear and bulk relaxation behaviors, the latter of which is absent in incompressible assumption of the continuum cell. Thus, by taking the compressibility and bulk relaxation into account, the optimized values for viscoelastic parameters were altered. Finally, by applying the NHPH model, the fluid flow was shown to contribute to the nonlinear and time-dependent creep of the cell inside the micropipette, which is represented by bulk relaxation in the NHVH model.

The optimized material parameters for incompressible NHVH model were in good agreement with those provided by Zhou et al. [9]. However, based on the present findings, including the compressibility in the model considerably

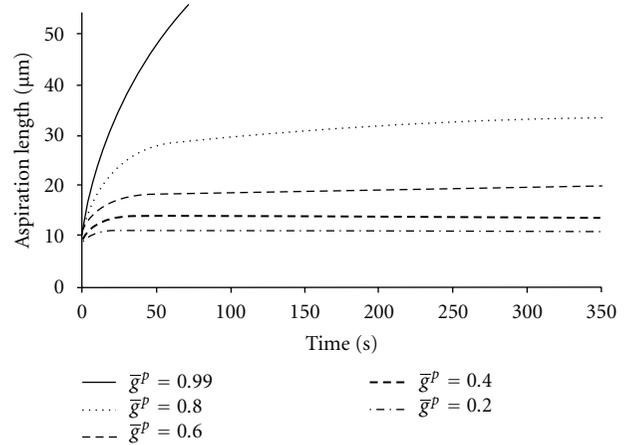


FIGURE 5: The FE results for NHVH cell model with different values of the Prony shear relaxation parameter \bar{g}^P . Other parameters were Prony bulk relaxation parameter $\bar{k}^P = 0$, time constant $\tau = 3.6$ (s), initial shear modulus $G_0 = 414$ (Pa), Poisson's ratio $\nu = 0.42$, and aspiration pressure of $P = 890$ (Pa).

enhanced the match between the FE model and the short-term experimental creep data points. The suggested Poisson's ratio of 0.42 is interestingly close to the values estimated for chondrocytes by Trickey et al. [8], and Jones et al. [7] and for THP-1 cells studied by Lin et al. [23].

The incompressible NHVH model with the material parameters obtained by applying the Sato et al. [4] analytical solution to the experimental data [5] showed a significantly different creep response of the cell than the experimental curve. This mainly arises from strain hardening, finite cell to pipette diameter, and large deformation regimen that are not taken into account in this analytical solution.

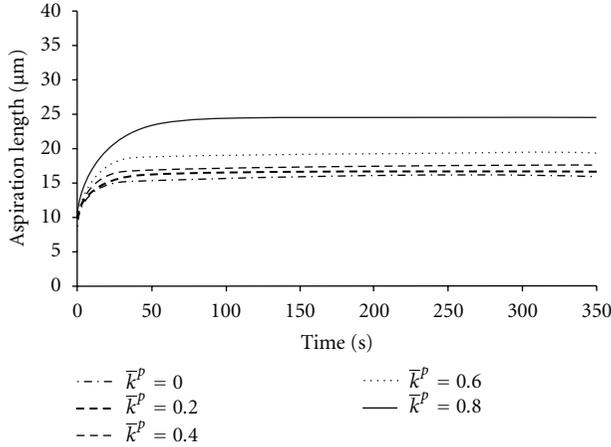


FIGURE 6: The FE results for NHHV cell model for different values of the Prony bulk relaxation parameter \bar{k}^P . Other parameters were $\bar{g}^P = 0.4$, time constant $\tau = 3.6$ (s), initial shear modulus $G_0 = 414$ (Pa), Poisson's ratio $\nu = 0.42$, and aspiration pressure $P = 890$ (Pa).

The final length of aspiration can be attributed to both bulk and shear relaxation of the cell material (Figures 5 and 6). Clearly, the bulk relaxation modulus is more influential with lower values of Poisson's ratio, while it has no effect with Poisson's ratio of 0.5 (Figure 7). The initial shear modulus obtained from the analytical solution was lower than what was estimated for the Neo-Hookean hyperelastic model. This leads to a difference in the viscoelastic parameters between the models. In the initial phase of aspiration, the length of protrusion increases almost linearly by decreasing Poisson's ratio. By assuming the cell as an incompressible material, there is only the shear behavior to enable the deformation of the cell, and hence a lower value for this parameter is estimated as compared to the compressible cell. This also leads to a higher value estimated for dimensionless shear modulus.

For comparison between different hyperelastic models, Arruda-Boyce viscohyperelastic (ABVH) material model was also tested. In agreement with the NHHV model, the compressible ABVH model could simulate well the experimental aspiration curve. The incompressible ABVH material model with a range of values presented in the cell and tissue mechanics literature [15, 34–36] could not reach the specified length of aspiration.

In the incompressible NHHV model, there were three material parameters to be optimized; G_0 , \bar{g}^P , and τ . Thus, the parameters of this model could be uniquely determined by three distinct regions in the aspiration length-time curve of the cell: the initial elastic jump of the cell inside the micropipette, the equilibrium length of aspiration, and the slope of the creep curve. In the compressible NHHV model, there were two additional parameters, D_1 (reflects Poisson's ratio and the initial shear modulus) and \bar{k}^P , which are mutually influencing each other (Figure 7), and hence they affect the uniqueness of the solution in this model. By assumption of a time-independent Poisson's ratio ($\bar{g}^P = \bar{k}^P$),

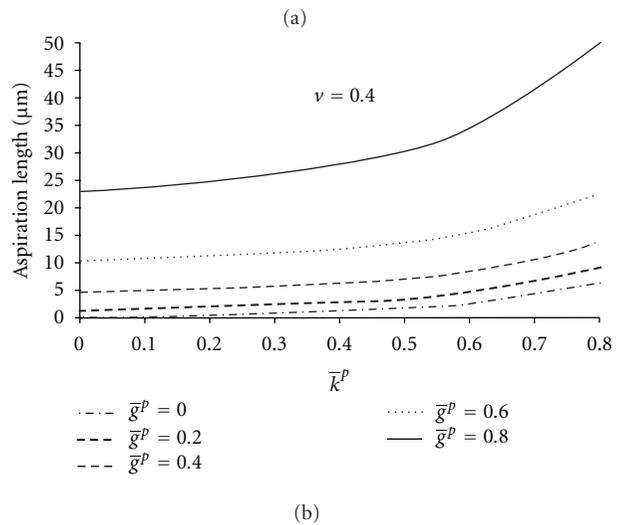
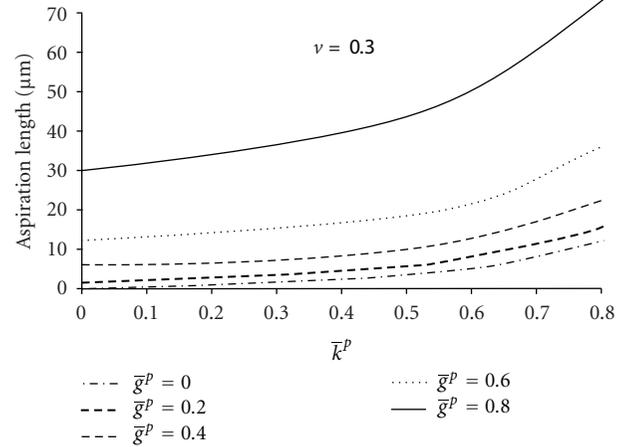


FIGURE 7: The NHHV FE model results for different values of Prony shear relaxation parameter, \bar{g}^P , and Prony bulk relaxation parameter, \bar{k}^P , with different Poisson's ratios: (a) $\nu = 0.3$ and (b) $\nu = 0.4$. Other parameters were time constant $\tau = 3.6$ (s), initial shear modulus $G_0 = 414$ (Pa), and aspiration pressure $P = 890$ (Pa). In this figure, the initial elastic jumps have been removed.

it was possible to uniquely calculate Poisson's ratio and the initial shear modulus, as \bar{g}^P was already known from the equilibrium time. However, this might not be fully realistic for viscoelastic materials in which Poisson's ratio is thought to vary with time. Even in that case, however, the optimum parameter for Poisson's ratio was always below 0.49.

The deformation-based bulk relaxation behavior may be an equivalent for the fluid flow gradient within the spongy cytoplasm [37]. In the present study, the NHPH model with a Poisson's ratio of 0.42 and permeability of less than 1.02×10^{-15} (m^4/Ns) showed a 17% aspiration length change between the initial elastic and poroelastic cell elongations. This is close to the amount of contribution of bulk relaxation creep in the aspiration curve (18.5% with the optimized compressible NHHV parameters reported in Table 3 and aspiration pressure of 445 Pa). Furthermore, with a lower Poisson's ratio the effect of fluid flow was amplified. During the first seconds in MA, the NHPH model behaves as an

incompressible material and the fluid flow is negligible. As a function of time during the creep, the fluid flows into lower-pressure areas. Depending on the compressibility, a large part of creep deformation was shown to be related to the volume relaxation due to fluid exchange. This result suggests that although the solid viscoelasticity, which mainly emerges from the cytoskeleton, has earlier been suggested to be the main contributor in the time-dependent cell deformation, the role of fluid may not be negligible.

Consistent with the present study, an instantaneous volume decrease was observed recently for chondrocytes in the *in vivo* loading of rabbit joints [38]. Furthermore, the compressibility was observed to be determinant in characterization of cell material properties and the pattern of tension distribution within the cytoplasm in compression of a cell by microplates [16]. This is in contrast with some observations in atomic force microscopy (AFM) and nanoindentation of cells where the effect of compressibility has been suggested to be insignificant [39]. Hence, the shear moduli of the cell measured with the MA technique are expected to be lower compared to their equivalents from the AFM indentation technique. This prediction is in agreement with a recent comparative study on giant phospholipid vesicles with AFM and MA techniques [40] and supported by Darling et al. [39] on Young's modulus of different mesenchymal lineage cells studied by different experimental techniques. In our study, the incompressible NHVH cell in MA technique underestimated the instantaneous and equilibrium Young's moduli by 20% and 41%, respectively, compared to the outcome of the optimized parameters for the compressible NHVH model. This is also in partial agreement with previous AFM indentation and MA of chondrocytes [6] where MA was observed to estimate 26% lower values for the equilibrium Young's modulus compared to AFM indentation technique. The apparent difference between the importance of compressibility and bulk relaxation in aspiration and elastic indentation of cells may originate from the different deformation directions the cell undergoes in each experiment and the different intracellular components involved as the load bearing components. In micropipette aspiration, the cell is mainly stretched, while in AFM test it is compressed. Based on the Tensegrity model for the cell [29] which is supported by later studies [41–43], the load bearing elements under pressure are thought to be microtubules while actin filaments bear the tensile load. The difference in mechanical properties of these subcellular elements may lead to diverging outcomes from different methods of inducing the deformation. From another point of view, the disparity may also originate from the different length and time scales of deformations in each experiment. In AFM indentation of the cells, the deformations are usually relatively small whereas in the micropipette aspiration the cell may experience large deformations/stretches. The time span of the experiment may also contribute to the discrepancy of outcomes between different techniques.

Due to the heterogeneity of the cells and anisotropic properties of the cells and membranes with compression-tension nonlinearity, the mechanical models used to interpret the results also contribute to discrepancies between the

outcomes of different techniques. Evaluating the capability of material models to mimic the cell behavior in different experimental conditions can facilitate a better understanding of the mechanical behavior of cells. Eventually, the models to describe the cell behavior realistically should give the same results in all testing geometries.

The cell membrane has a very low permeability [31, 32]. In the present study, the membrane was not modeled as a separate layer, but the effect of the membrane and cell permeability on creep deformation was simulated by implementing the minimum and maximum values reported for the cell [30] and membrane [31, 32] in the homogenized cell model. Even though considering the cell membrane as a separate layer with different permeability than that of the cell would be favorable to obtain a detailed distribution of stress, strain, and fluid pressure forming within the cell and to examine pressure driven fluid flow across the membrane, the homogenized cell permeability was considered enough here.

The current model was based on a continuum assumption, and detailed intracellular elements were not included, allowing us to obtain unambiguous values of material parameters for the cell. In order to provide a detailed and realistic model of the cell that can explain the chemo-mechanical coupling of the cell-ECM interactions, the compressibility of the cytoskeleton and the contribution of the nucleus and cell membrane to the overall mechanical properties of the cell should be considered in the future.

5. Conclusions

Parametric studies and optimization results suggest that compressibility and bulk relaxation behavior are two important factors in the deformation behavior of the cell in MA technique, which are not considered in the commonly used equation for MA. The compressibility of the cell that is presented by Poisson's ratio might explain the behavior of microstructural network of the cytoskeleton. Intracytoplasmic and transmembrane fluid flow could be responsible for bulk relaxation behavior and volume change of cells in MA. These mechanisms, if present in cells *in vivo*, may modulate significantly cell mechanics and mechanotransduction in biological tissues. Further, poroelastic fluid flow and compressibility of cells in multiscale models of biological tissues may change the prediction of cell responses *in situ/in vivo*.

Acknowledgments

Financial support from the Academy of Finland (projects 125415, 140730, and 218038), Sigrid Jusélius Foundation, Finland, European Research Council (ERC), and National Doctoral Programme of Musculoskeletal Disorders and Biomaterials, Finland, is acknowledged. CSC-IT Center for Science, Finland is acknowledged for technical support. Finally, we would like to thank Professor Mahdi Navidbakhsh, Iran University of Science and Technology, for his consultation.

References

- [1] C. T. Lim, E. H. Zhou, A. Li, S. R. K. Vedula, and H. X. Fu, "Experimental techniques for single cell and single molecule biomechanics," *Materials Science and Engineering C*, vol. 26, no. 8, pp. 1278–1288, 2006, Proceedings of the 1st TMS Symposium on Biological Materials Science.
- [2] G. Bao and S. Suresh, "Cell and molecular mechanics of biological materials," *Nature Materials*, vol. 2, no. 11, pp. 715–725, 2003.
- [3] I. Titushkin and M. Cho, "Modulation of cellular mechanics during osteogenic differentiation of human mesenchymal stem cells," *Biophysical Journal*, vol. 93, no. 10, pp. 3693–3702, 2007.
- [4] M. Sato, D. P. Theret, L. T. Wheeler, N. Ohshima, and R. M. Nerem, "Application of the micropipette technique to the measurement of cultured porcine aortic endothelial cell viscoelastic properties," *Journal of Biomechanical Engineering*, vol. 112, no. 3, pp. 263–268, 1990.
- [5] S. C. W. Tan, W. X. Pan, G. Ma, N. Cai, K. W. Leong, and K. Liao, "Viscoelastic behaviour of human mesenchymal stem cells," *BMC Cell Biology*, vol. 9, p. 40, 2008.
- [6] E. M. Darling, S. Zauscher, and F. Guilak, "Viscoelastic properties of zonal articular chondrocytes measured by atomic force microscopy," *Osteoarthritis and Cartilage*, vol. 14, no. 6, pp. 571–579, 2006.
- [7] W. R. Jones, H. P. Ting-Beall, G. M. Lee, S. S. Kelley, R. M. Hochmuth, and F. Guilak, "Alterations in the young's modulus and volumetric properties of chondrocytes isolated from normal and osteoarthritic human cartilage," *Journal of Biomechanics*, vol. 32, no. 2, pp. 119–127, 1999.
- [8] W. R. Trickey, F. P. T. Baaijens, T. A. Laursen, L. G. Alexopoulos, and F. Guilak, "Determination of the Poisson's ratio of the cell: recovery properties of chondrocytes after release from complete micropipette aspiration," *Journal of Biomechanics*, vol. 39, no. 1, pp. 78–87, 2006.
- [9] E. H. Zhou, C. T. Lim, and S. T. Quek, "Finite element simulation of the micropipette aspiration of a living cell undergoing large viscoelastic deformation," *Mechanics of Advanced Materials and Structures*, vol. 12, no. 6, pp. 501–512, 2005.
- [10] J. Xu, Y. Tseng, and D. Wirtz, "Strain hardening of actin filament networks: regulation by the dynamic cross-linking protein α -actinin," *Journal of Biological Chemistry*, vol. 275, no. 46, pp. 35886–35892, 2000.
- [11] D. Stamenović and D. E. Ingber, "Models of cytoskeletal mechanics of adherent cells," *Biomechanics and Modeling in Mechanobiology*, vol. 1, no. 1, pp. 95–108, 2002.
- [12] J. Ohayon and P. Tracqui, "Computation of adherent cell elasticity for critical cell-bead geometry in magnetic twisting experiments," *Annals of Biomedical Engineering*, vol. 33, no. 2, pp. 131–141, 2005.
- [13] N. Caille, O. Thoumine, Y. Tardy, and J. J. Meister, "Contribution of the nucleus to the mechanical properties of endothelial cells," *Journal of Biomechanics*, vol. 35, no. 2, pp. 177–187, 2002.
- [14] P. Tracqui, J. Ohayon, and T. Boudou, "Theoretical analysis of the adaptive contractile behaviour of a single cardiomyocyte cultured on elastic substrates with varying stiffness," *Journal of Theoretical Biology*, vol. 255, no. 1, pp. 92–105, 2008.
- [15] I. Kang, D. Panneerselvam, V. P. Panoskaltis, S. J. Eppell, R. E. Marchant, and C. M. Doerschuk, "Changes in the hyperelastic properties of endothelial cells induced by tumor necrosis factor- α ," *Biophysical Journal*, vol. 94, no. 8, pp. 3273–3285, 2008.
- [16] J. P. McGarry, "Characterization of cell mechanical properties by computational modeling of parallel plate compression," *Annals of Biomedical Engineering*, vol. 37, no. 11, pp. 2317–2325, 2009.
- [17] V. Lulevich, T. Zink, H. Y. Chen, F. T. Liu, and G. Y. Liu, "Cell mechanics using atomic force microscopy-based single-cell compression," *Langmuir*, vol. 22, no. 19, pp. 8151–8155, 2006.
- [18] K. B. Bernick, T. P. Prevost, S. Suresh, and S. Socrate, "Biomechanics of single cortical neurons," *Acta Biomaterialia*, vol. 7, no. 3, pp. 1210–1219, 2011.
- [19] G. U. Unnikrishnan, V. U. Unnikrishnan, and J. N. Reddy, "Constitutive material modeling of cell: a micromechanics approach," *Journal of Biomechanical Engineering*, vol. 129, no. 3, pp. 315–323, 2007.
- [20] M. A. Haider and F. Guilak, "Application of a three-dimensional poroelastic BEM to modeling the biphasic mechanics of cell-matrix interactions in articular cartilage," *Computer Methods in Applied Mechanics and Engineering*, vol. 196, no. 31–32, pp. 2999–3010, 2007.
- [21] D. Zhang, "Oscillatory pressurization of an animal cell as a poroelastic spherical body," *Annals of Biomedical Engineering*, vol. 33, no. 9, pp. 1249–1269, 2005.
- [22] E. Kim, F. Guilak, and M. A. Haider, "The dynamic mechanical environment of the chondrocyte: a biphasic finite element model of cell-matrix interactions under cyclic compressive loading," *Journal of Biomechanical Engineering*, vol. 130, no. 6, Article ID 061009, 2008.
- [23] L. A. G. Lin, A. Q. Liu, Y. F. Yu et al., "Cell compressibility studies utilizing noncontact hydrostatic pressure measurements on single living cells in a microchamber," *Applied Physics Letters*, vol. 92, no. 23, Article ID 233901, 3 pages, 2008.
- [24] Dassault Systemes Simulia Corp. Abaqus User's Manual, v. 6.8.1. Providence, RI, USA, 2008.
- [25] J. Yan and J. S. Strenkowski, "A finite element analysis of orthogonal rubber cutting," *Journal of Materials Processing Technology*, vol. 174, no. 1–3, pp. 102–108, 2006.
- [26] P. Julkunen, W. Wilson, J. S. Jurvelin, and R. K. Korhonen, "Composition of the pericellular matrix modulates the deformation behaviour of chondrocytes in articular cartilage under static loading," *Medical and Biological Engineering and Computing*, vol. 47, no. 12, pp. 1281–1290, 2009.
- [27] M. A. Haider and F. Guilak, "An axisymmetric boundary integral model for assessing elastic cell properties in the micropipette aspiration contact problem," *Journal of Biomechanical Engineering*, vol. 124, no. 5, pp. 586–595, 2002.
- [28] D. Banabic, M. Bambach, M. Cannamela, M. Azaouzi, G. Hirt, and J. L. Batoz, "Computer-aided tool path optimization for single point incremental sheet forming," in *Advanced Methods in Material Forming*, p. 362, 2007.
- [29] A. J. Maniatis, C. S. Chen, and D. E. Ingber, "Demonstration of mechanical connections between integrins, cytoskeletal filaments, and nucleoplasm that stabilize nuclear structure," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 94, no. 3, pp. 849–854, 1997.
- [30] G. T. Charras, P. P. Lehenkari, and M. A. Horton, "Atomic force microscopy can be used to mechanically stimulate osteoblasts and evaluate cellular strain distributions," *Ultra-microscopy*, vol. 86, no. 1–2, pp. 85–95, 2001.
- [31] J. Soulhat, M. D. Buschmann, and A. Shirazi-Adl, "A fibril-network-reinforced biphasic model of cartilage in unconfined

- compression,” *Journal of Biomechanical Engineering*, vol. 121, no. 3, pp. 340–347, 1999.
- [32] R. Krishnan, S. Park, F. Eckstein, and G. A. Ateshian, “Inhomogeneous cartilage properties enhance superficial interstitial fluid support and frictional properties, but do not provide a homogeneous state of stress,” *Journal of Biomechanical Engineering*, vol. 125, no. 5, pp. 569–577, 2003.
- [33] V. C. Mow, S. C. Kuei, W. M. Lai, and C. G. Armstrong, “Biphasic creep and stress relaxation of articular cartilage in compression: theory and experiments,” *Journal of Biomechanical Engineering*, vol. 102, no. 1, pp. 73–84, 1980.
- [34] Y. Liu, A. E. Kerdok, and R. D. Howe, “A nonlinear finite element model of soft tissue indentation,” in *Proceedings of the Medical Simulation: International Symposium (ISMS '04)*, pp. 67–76, 2004.
- [35] T. P. Prevost, A. Balakrishnan, S. Suresh, and S. Socrate, “Biomechanics of brain tissue,” *Acta Biomaterialia*, vol. 7, no. 1, pp. 83–95, 2011.
- [36] A. Balakrishnan and S. Socrate, “Material property differentiation in indentation testing using secondary sensors,” *Experimental Mechanics*, vol. 48, no. 4, pp. 549–558, 2008.
- [37] G. T. Charras, T. J. Mitchison, and L. Mahadevan, “Animal cell hydraulics,” *Journal of Cell Science*, vol. 122, no. 18, pp. 3233–3241, 2009.
- [38] Z. Abusara, R. Seerattan, A. Leumann, R. Thompson, and W. Herzog, “A novel method for determining articular cartilage chondrocyte mechanics in vivo,” *Journal of Biomechanics*, vol. 44, no. 5, pp. 930–934, 2011.
- [39] E. M. Darling, M. Topel, S. Zauscher, T. P. Vail, and F. Guilak, “Viscoelastic properties of human mesenchymally-derived stem cells and primary osteoblasts, chondrocytes, and adipocytes,” *Journal of Biomechanics*, vol. 41, no. 2, pp. 454–464, 2008.
- [40] S. Dieluweit, A. Csiszár, W. Rubner, J. Fleischhauer, S. Houben, and R. Merkel, “Mechanical properties of bare and protein-coated giant unilamellar phospholipid vesicles. A comparative study of micropipet aspiration and atomic force microscopy,” *Langmuir*, vol. 26, no. 13, pp. 11041–11049, 2010.
- [41] T. J. Dennerll, H. C. Joshi, V. L. Steel, R. E. Buxbaum, and S. R. Heidemann, “Tension and compression in the cytoskeleton of PC-12 neurites II: quantitative measurements,” *Journal of Cell Biology*, vol. 107, no. 2, pp. 665–674, 1988.
- [42] N. Wang, K. Naruse, D. Stamenović et al., “Mechanical behavior in living cells consistent with the tensegrity model,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 98, no. 14, pp. 7765–7770, 2001.
- [43] D. Stamenović, N. Wang, and D. E. Ingber, “Tensegrity architecture and the mammalian cell cytoskeleton,” in *Multiscale in Molecular and Continuum Mechanics: Interaction of Time and Size from Macro to Nano*, pp. 321–338, 2007.

Research Article

Higher-Order Spectrum in Understanding Nonlinearity in EEG Rhythms

Cauchy Pradhan,¹ Susant K. Jena,² Sreenivasan R. Nadar,³ and N. Pradhan⁴

¹Medical Image Processing Lab, EPFL, 1015 Lausanne, Switzerland

²Department of Computer Science and Mathematics, K.N.S. Institute of Technology, Bangalore 560064, India

³MEG Core, National Institute of Mental Health, Bethesda, MD 20892, USA

⁴Department of Psychopharmacology, NIMHANS, Bangalore 560029, India

Correspondence should be addressed to N. Pradhan, nprnimhans@gmail.com

Received 2 October 2011; Accepted 27 October 2011

Academic Editor: Vikas Rai

Copyright © 2012 Cauchy Pradhan et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

The fundamental nature of the brain's electrical activities recorded as electroencephalogram (EEG) remains unknown. Linear stochastic models and spectral estimates are the most common methods for the analysis of EEG because of their robustness, simplicity of interpretation, and apparent association with rhythmic behavioral patterns in nature. In this paper, we extend the use of higher-order spectrum in order to indicate the hidden characteristics of EEG signals that simply do not arise from random processes. The higher-order spectrum is an extension Fourier spectrum that uses higher moments for spectral estimates. This essentially nullifies all Gaussian random effects, therefore, can reveal non-Gaussian and nonlinear characteristics in the complex patterns of EEG time series. The paper demonstrates the distinguishing features of bispectral analysis for chaotic systems, filtered noises, and normal background EEG activity. The bispectrum analysis detects nonlinear interactions; however, it does not quantify the coupling strength. The squared bicoherence in the nonredundant region has been estimated to demonstrate nonlinear coupling. The bicoherence values are minimal for white Gaussian noises (WGNs) and filtered noises. Higher bicoherence values in chaotic time series and normal background EEG activities are indicative of nonlinear coupling in these systems. The paper shows utility of bispectral methods as an analytical tool in understanding neural process underlying human EEG patterns.

1. Introduction

Biological signals are highly complex and understandably nonlinear, may it be the firing of neurons, the beating of the heart, or breathing. The nature of the signals is dramatic and appears to be esoteric. These signals arise out of a multitude of interconnected elements comprising of the human body. These are bounded, finite, and the connections are weakly coupled across all elements. These vary in time scales ranging from nanoseconds for molecular motion to gross observable behavior in terms of days and years. This implies that biological systems are nonstationary [1, 2]. Signals of biological origin require considerable experience for their analysis and interpretation. This is only gained through practical training and hand on expertise as most signals are contaminated by processes of nonbiological origin at the

recording stage. Further, these signals can be mimicked by noise and artifacts [3]. One of the contentious issues is that filtered noises mimick time series of brain electrical potentials which substantially vitiate both extraction of linear and nonlinear measures of these time series.

Electroencephalography (EEG) is the recording of brain cortical electrical activity from electrodes placed on the scalp. The signals are also recorded subdurally or directly on the cortex and are called Electrocochogram (ECoG). The dynamics and the structure patterns of both EEG and ECoG are similar. EEG signals are used in seizure detection, organic encephalopathies, monitoring anesthesia, and for the determination of brain death. There are perhaps 10^5 neurons under each square millimeter of the cortical surface. Scalp EEG measures space-averaged activity of 10^7 or more neurons implying a source area of at least a square centimeter. EEG

signals are quasistationary not exceeding 2–4 seconds in eye-open state or 20–30 seconds in deep sleep. The stationary condition varies with fleeting attention, and thus analytical methods based on the assumption of the stationarity of the system are superfluous. Signal prediction, therefore, is fallacious other than the mere detection of trends. Under these circumstances, linear stochastic models remain the primary method for the analysis of biological time series due to their robustness, simplicity of interpretation, and their apparent association with rhythmic behavioral patterns in nature. The Fourier spectrum is trivial to linear methods. The linear algebraic techniques and spectral estimates over years have shown good correlation as approximation for inherently nonlinear functions that are biological data [4–7].

Over decades of research publications, the application of nonlinear time series methods involving invariant measures such as Correlation Dimension (D_2), Lyapunov Exponents (λ), and Fractal Dimension (D_1), have not gained wide acceptance in biomedical field. The measures do not provide real insight into the biological phenomena due to their inherent high dimensionality along with long- and short-range interactions within these systems [8–12]. The short data lengths of biological time series deny perfect reconstruction of their attractor. In attractor reconstruction, true independence of each vector is not guaranteed due to the partial correlations across dimensions. Therefore, nonlinear invariant measures like D_2 , D_1 , and λ have limited value in medicine. Even if one establishes the presence of chaos in deep sleep or anesthesia [8–12], it is very difficult to reconstruct a model that can reproduce the time series reflecting the behavioral transitions during sleep and their apparent associations with the time series. More so, the filtration of WGN using stiff linear filters can mimic chaotic process including its validity by surrogate testing [3, 13–16]. The results of nonlinear time series invariants could not be validated as two different persons in the same behavioral states have varying values. We have reported D_2 values in the range of 2–4 in seizures and deep sleep. Similar D_2 values are also found in normal healthy individuals in awake state [17–22]. Therefore, the biological acceptability of published nonlinear time series measures is low. Even if we get a long-range pattern, there remain high degrees of overlap in such invariant measures across behavioral states.

In order to elicit better inference of EEG time series data, we have extended the Fourier transform to bispectral estimation using higher-order (third) moment [23–26]. This nullifies all Gaussian random effects in the process. While bispectrum analysis detects nonlinear interactions, it does not quantify the coupling strength which we have evaluated through bicoherence. The bicoherence, that is, a normalized bispectrum has been used in sleep studies in animals [27]. There are also several reports of its application to human brain signals, mainly, monitoring the degree of consciousness or depth of anaesthesia. Bispectral methods have led to the development of a device (Bispectral Index Monitor) which has been recommended for clinical monitoring in critical care and surgical anesthesia [28]. Recently, bispectral measures have been extended to detect subtle changes in EEG dynamics in visual representation of motor tasks [29].

However, application of higher-order spectral techniques is less wide-spread in medicine and biology. The aim of the present investigation is to demonstrate the utilities of higher-order spectrum in human EEG processing. There is a need to represent the standard rhythmic manifestations of EEG as revealed by bispectrum and reinforce its clinical and research applications. All descriptions of brain electrical activity recordings invariably refer to these well-established background EEG rhythms. This paper also demonstrates the significance of nonlinear interactions and coupling of brain activity in the cerebral cortex using background EEG activities of normal subjects. We have used delta, theta, alpha, beta, and indeterminate activities for the above study and contrasted the results with chaotic time series (Lorenz and Mackey-Glass systems), white Gaussian noise, and filtered noises (1.5 Hz, 3 Hz, 9 Hz, 30 Hz, and 300 Hz).

2. Method

Any process is a linear process with respect to its second-order statistics (power spectrum). The autocorrelation sequence does not give any evidence of nonlinearity. In contrast, higher-order cumulants can give evidence of nonlinearity (i.e., bispectrum for quadratic interactions). Polyspectrum estimators are the natural generalizations of the autocorrelation function, and cumulants are specific nonlinear combinations of these moments. The power spectrum does not carry information about phase which can be recovered from higher-order polyspectra. The use of higher-order moments nullifies all Gaussian random effects of the process, and the bicoherence can then quantify the degree of the remaining nonlinear coupling. The bispectrum and its normalized derivative, the bicoherence, describes the components of a time series that deviates from a Gaussian amplitude distribution. Bispectra have been used to examine various physical time series data including plasma physics and ocean waves [23–26].

Fourier transform of $f(t)$ is given by

$$f'(w) = \int_{-\infty}^{\infty} f(t)e^{-i\omega t} dt. \quad (1)$$

Power spectrum of (1) is

$$Pf(w) = |f'(w)|^2. \quad (2)$$

The natural estimate of the bispectrum (B_{xxx}) is the Fourier transform of the third-order cumulant sequence (C_{xxx}):

$$B_{xxx}(f_1, f_2) = \sum_{k=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} C_{xxx}(k, l) e^{-i2\pi f_1 k} e^{-i2\pi f_2 l}. \quad (3)$$

The bispectrum can be also written as (4), where X_n is the Fourier Transform of $\{X_n\}$:

$$B_{xxx}(f_1, f_2) = \frac{1}{N^2} X_n'(f_1 + f_2) X_n f_1 X_n f_2. \quad (4)$$

The bicoherence or the normalized bispectrum (5) is a measure of the amount of phase coupling that occurs in a

signal or between two signals. Phase coupling is the estimate of the proportion of energy in every possible pair of frequency components, f_1, f_2 (i.e., 1–50 Hz in EEG), which satisfies the definition of quadratic phase coupling (phase of component at f_3 , which is $f_1 + f_2$, equals phase of f_1 + phase of f_2) [23, 25]:

$$\text{bic}(f_1, f_2) = \frac{|B(f_1, f_2)|^2}{P(f_1)P(f_2)P(f_1 + f_2)}. \quad (5)$$

When the analyzed signal exhibits structure of any kind whatsoever, it might be expected that some phase coupling occurs. Bicoherence analysis is able to detect coherent signals in extremely noisy data, provided that the coherency remains constant for sufficiently long times, since the noise contribution falls off rapidly with increasing N . The bicoherence due to the coherent signal should be at least a factor of 3 above the maximum noise contribution, and the allowable signal-to-noise ratio is 3e. Bicoherence can be considered a very powerful noise filter. We have utilized the archived EEG data available in the Department of Psychopharmacology at the National Institute of Mental Health and Neurosciences, Bangalore. We have taken normative data of standard frequency band of EEG patterns of alpha, beta, theta, delta and indeterminate activities. These pattern descriptions are the accepted background rhythms in the EEG literature.

3. Results

The Lorenz and the Mackey-Glass systems are examples of classical nonlinear and chaotic systems. These are included in the study for the purpose of comparison with linear systems (WGN, filtered noises) and EEG with a view to demonstrate non-Gaussian and nonlinear characteristics of EEG. The time series of the Lorenz (x -component) and the Mackey-Glass systems are shown in Figure 1. The WGN and filtered noises of 1.5 Hz, 3 Hz, 9 Hz, 30 Hz, and 300 Hz used in the study are given in Figure 2. The EEG signals of alpha, beta, theta, delta, and indeterminate activities are shown in Figure 3. We have used 4096 data points of the above time series for estimation of bicoherence and Hinich statistics. The respective bicoherence plots are given in Figures 4–6. The numerical simulation of chaotic time series, WGN, filtered noises, and the method of EEG data acquisition are described in [25].

Table 1 provides probability measure of significance (χ^2 and P values), where the assumption Gaussianity holds good for $P \geq 0.5$. The results show $P = 0$ for chaotic time series, low-frequency filtered noises (1.5 Hz–30 Hz), and EEG data. Therefore, non-Gaussianity holds good for these time series. The low-frequency filtered noises (1.5 Hz to 30 Hz) are within the bandwidth of normal background EEG activities; therefore, the filtered noises may be considered as spurious processes mimicking EEG rhythms. The results given in Table 1 indicate large differences in the estimated and theoretical interquartile ranges for chaotic time series and the EEG data. In contrast, the estimated and theoretical interquartile ranges (R) are relatively small for the WGN and filter noises. As there are no gross differences between them,

TABLE 1

Time series	χ^2	P	R - estimated	R - theoretical	λ
Lorenz system	8469.61	0.0	589.87	46.87	301.35
Mackey-Glass system	5765.88	0.00	408.36	38.49	203.02
White Gaussian noise	2.07	1.00	0.08	2.25	0.07
Filtered noise 1.5 Hz	270.08	0.00	16.31	8.67	9.93
Filtered noise 3 Hz	627.06	0.00	46.60	12.99	22.74
Filtered noise 9 Hz	359.46	0.00	19.83	9.95	13.20
Filtered noise 30.0 Hz	120.05	0.00	5.80	5.89	4.40
Filtered noise 300 Hz	55.68	0.21	0.61	4.03	1.84
Alpha EEG activity	9731.75	0.00	390.91	35.75	175.05
Beta EEG activity	34794.08	0.00	1407.53	67.65	628.16
Theta EEG activity	4442.43	0.00	118.94	23.49	75.33
Delta EEG activity	14962.96	0.00	602.83	44.29	269.02
Indeterminate EEG activity	3178.88	0.00	114.14	20.06	54.83

χ^2 value, P : probability of the time series being Gaussian, R : interquartile range, λ : noncentrality parameter, *d.o.f* = 271.

the assumptions of linearity are not rejected in these data. The large differences in R -estimated and R -theoretical indicate nonlinearity. Therefore, the normal background EEG rhythms may be treated as nonlinear like chaotic processes. The results clearly establish non-Gaussian and nonlinear nature of background EEG activities.

4. Discussion

Higher-order statistics (spectra) and their application to various signal processing problems are relatively recent. There may be much more information in a stochastic non-Gaussian or deterministic signal than conveyed by its autocorrelation or power-spectral estimates. The higher-order spectra which is defined in terms of the higher-order moments or cumulants of a signal may contain this additional information. The higher-order moments are the natural generalizations of autocorrelation, and the cumulants are specific nonlinear combinations of these moments. The n th-order spectrum is defined as the Fourier transform of the n th-order cumulant sequence. The test for Gaussianity and linearity is based on the assumption that if the third-order cumulant of a process is zero, then its bispectrum is zero, and hence the bicoherence is also zero. A nonzero bispectrum, therefore, holds good for a non-Gaussian process [23, 30].

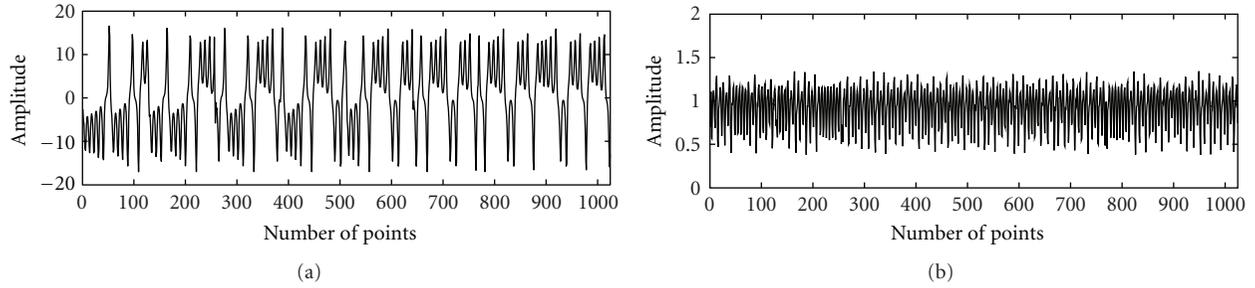


FIGURE 1: The time series of the (a) Lorenz (x -component) and (b) the Mackey-Glass systems. For each, 200000 data points were generated of which first 5000 data points were discarded to remove the initial transients. Only 1024 data points are shown here.

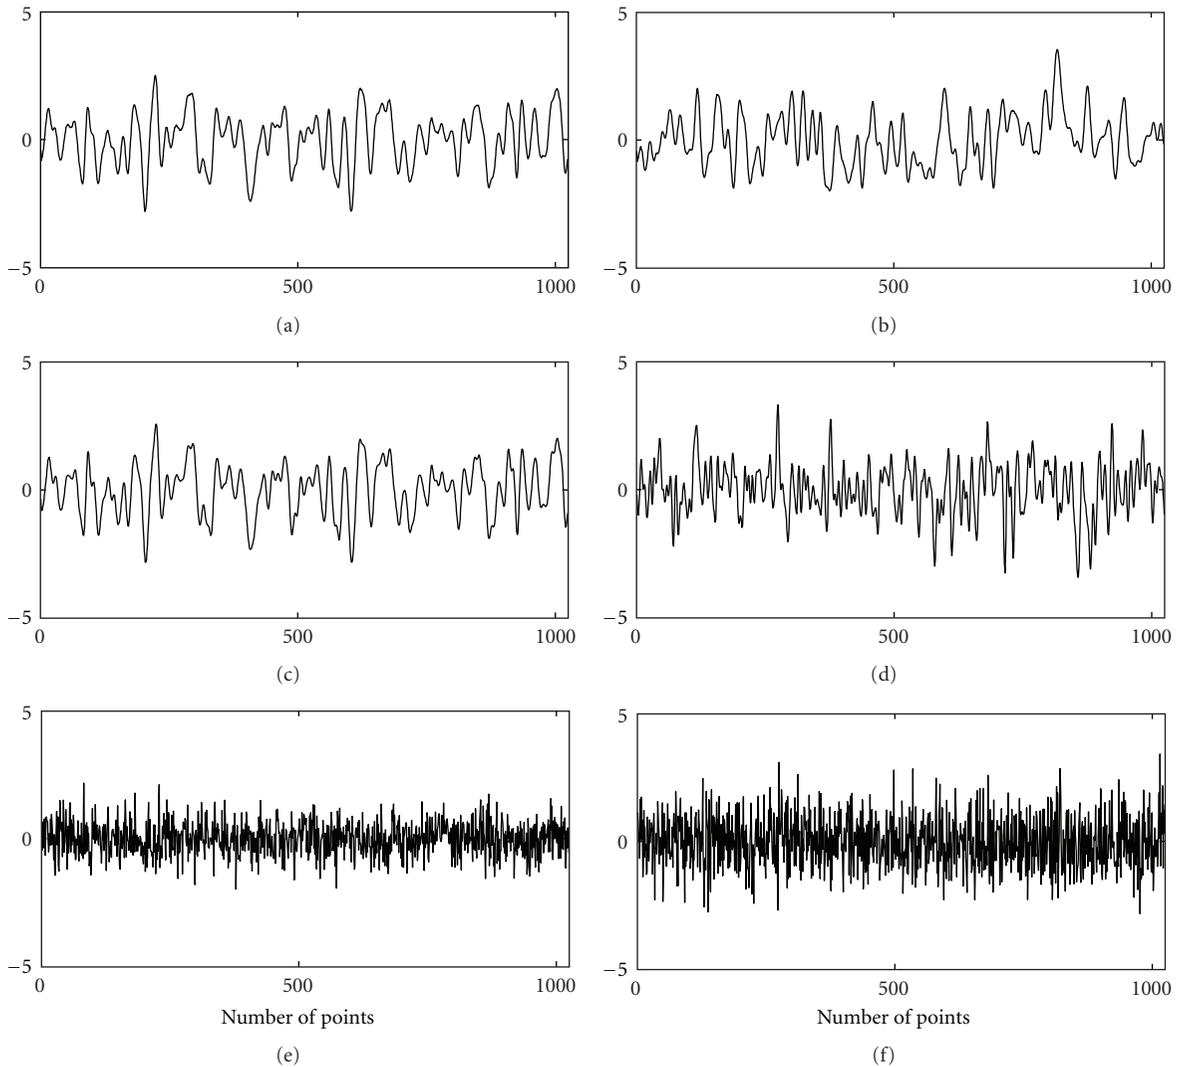


FIGURE 2: The plots show time series data of (a) 1.5 Hz low-pass filtered WGN, (b) 3 Hz low-pass filtered WGN, (c) 9 Hz low-pass filtered WGN, (d) 30 Hz low-pass filtered WGN, (e) 300 low-pass filtered WGN, and (f) White Gaussian noise.

For a linear process, the bicoherence is a nonzero constant. If the bispectrum is Gaussian distributed, we know that the squared bispectrum is chi-square distributed with two degrees of freedom. If the estimated bispectrum is zero, then the statistic of bicoherence is a central chi-square

random variable with two degrees of freedom. The squared bicoherence is summed over m points in the nonredundant region. Then resulting statistic is chi-square distributed with $2m$ degrees of freedom. The statistical test determines the consistency of bicoherence values with a central chi-square

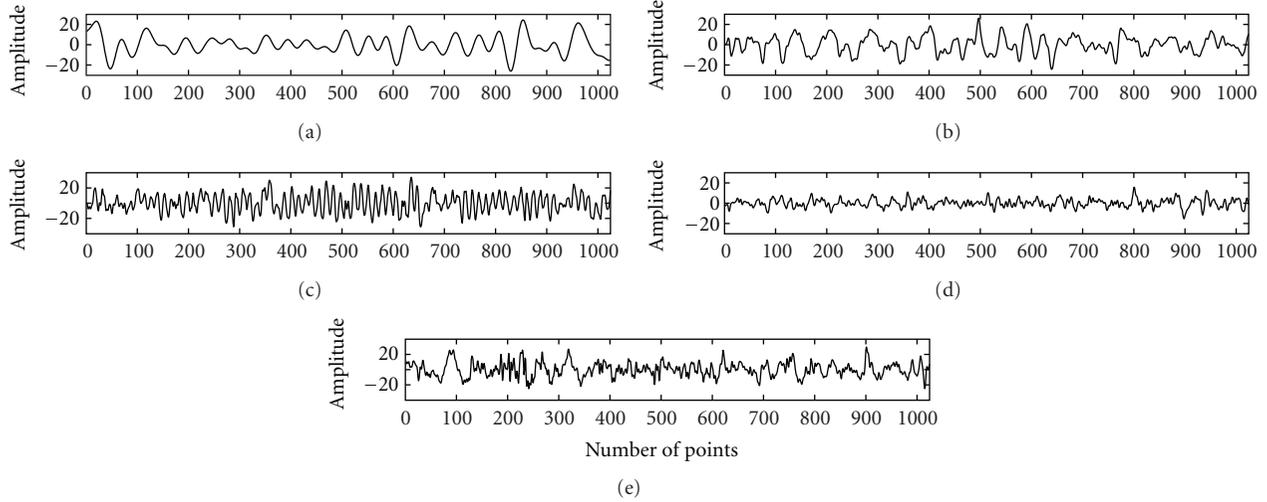


FIGURE 3: The plots show experimental time series data of various EEG activities: (a) delta EEG activity, (b) theta EEG activity, (c) alpha EEG activity, (d) beta EEG activity, and (e) indeterminate EEG activity.

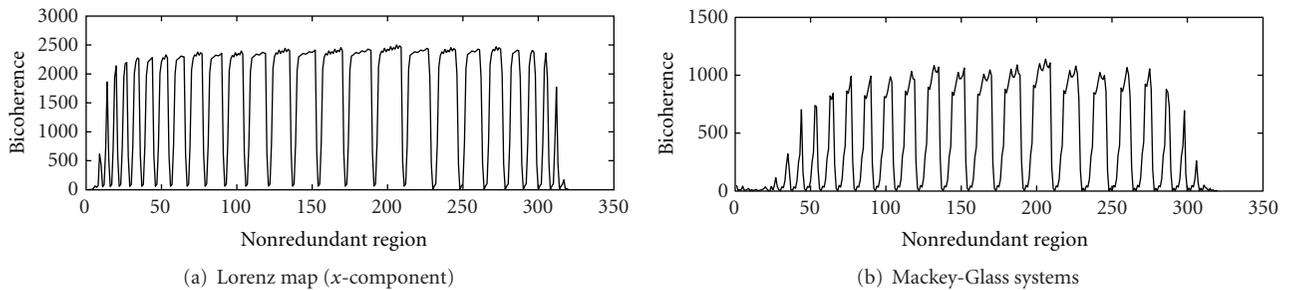


FIGURE 4: The bicoherence plots of classical chaotic systems: (a) Lorenz map (x -component), (b) Mackey-Glass systems.

distribution. This method is suitable for extracting information about the structure of the signals as it preserves nonminimum phase information. The higher-order spectra carry phase information which is ordinarily suppressed in power spectral estimate. Extensive studies in the field of signal processing have generated information on the input-output relationships of linear systems through autocorrelation and power spectrum. However, no definite information is available on the input-output relationship of nonlinear systems to stochastic excitation and each type of nonlinearity is treated as a special case [23, 30–32].

The present investigation uses the above concepts of higher-order spectral (bicoherence) estimates to obtain Hinich statistics for Gaussianity and linearity. The aim of the study is to contrast the Hinich statistics and bicoherence for well-understood classical chaotic system, WGN, filtered noises, and the normal background EEG activities. The results are shown in Table 1 which provides the χ^2 values, P values, R -estimated, R -theoretical, and λ values. These estimated values of EEG activities may be viewed in reference to obvious nonlinear chaotic series of Lorenz and Mackey-Glass systems and that of filtered noises which fall within its bandwidth.

The bispectral analyses of chaotic systems are found to be characteristically non-Gaussian and nonlinear. The chaos hypothesis in biological systems is based on a finite correlation dimension and a positive dominant Lyapunov exponent. A number of technical problems, however, confound the estimation of these measures. Even the instrumentation at the recording stage may render the application of nonlinear dynamics to these signals invalid. The use of analog filters during signal capture and digital filters in its analysis may account for spurious D_2 or positive dominant Lyapunov exponent values [25, 33], whereas bispectral methods are robust and least affected by extraneous noises and filters. Here (Table 1), the chaotic time series (Lorenz x -component and Mackey-Glass systems) are not only non-Gaussian but also have skewed noncentral distribution (λ). In addition, there is a large difference in their R -estimated and R -theoretical interquartile ranges. The Gaussianity assumption is rejected as the probability of the time series being Gaussian is small ($P = 0$). The linearity hypothesis also cannot be accepted since the differences in the estimated interquartile ranges are much larger than the theoretical values. The squared bicoherence in nonredundant regions are positive and have several peaks (Figure 4); hence, nonlinear coupling is indicated in this systems.

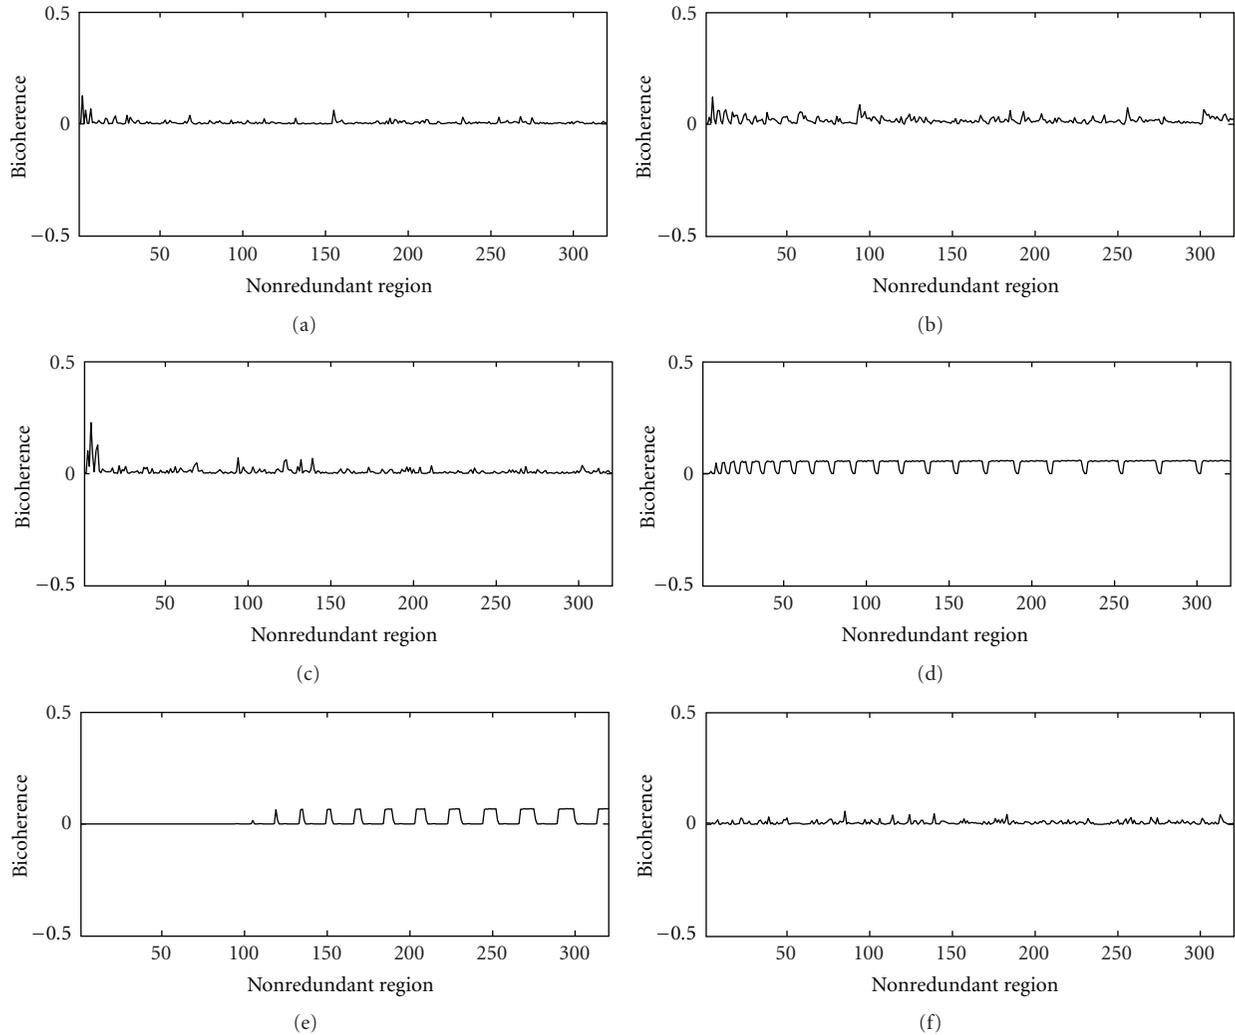


FIGURE 5: Bicoherence plots of (a) 1.5 low-pass filtered WGN, (b) 3 Hz low-pass filtered WGN, (c) 9 Hz low-pass filtered WGN, (d) 30 Hz low-pass filtered WGN, (e) 300 low-pass filtered WGN, and (f) White Gaussian noise.

The Hinich statistics (Table 1) also includes the results for WGN, filter noises. The WGN obviously shows highest probability ($P = 1$) as to being Gaussian. Therefore, the results of linearity test are ignored. The stiffly filtered noises (1.5 Hz, 3 Hz, 9 Hz, and 30 Hz) in the low-frequency ranges do depart from Gaussianity ($P = 0$); however, the estimated and theoretical interquartile ranges are very close to each other. The 300 Hz signal is more closer to WGN than the low-frequency signals. Therefore, the Gaussianity test holds good at this high frequency. The linearity hypothesis is also valid for WGN and filtered noises since their estimated and theoretical interquartile ranges (R) are close to one another. Figure 5 shows the estimated mean values of bicoherence over the points in the nonredundant region. The plots do not indicate nonlinear interactions or coupling.

The EEG time series have the lowest probability ($P = 0$) of being non-Gaussian with large differences in their estimated and theoretical interquartile ranges. The high non-central characteristics (λ) are seen only with chaotic and EEG

time series. The mean bicoherence of EEG signals are positive and large in the nonredundant regions with several peaks (Figure 6). The bicoherent values are found to be high for delta activity compared to other EEG rhythms. It reflects neocortical forcing for nonlinear coupling by the low-frequency neurons in the deep midbrain and the brain stem structures. We have shown higher bicoherence for all of the normal background EEG rhythms that are associated with a various behavioral states ranging from waking state and alert behavior (beta and indeterminate) to light sleep (alpha and theta) and deep sleep (delta). The normal background EEG activities have bicoherence values ranging from 90 for alpha to 4594.42 for delta. The high-squared bicoherence in nonredundant regions reflect quadratic phase coupling of neuronal ensembles in these conditions. The synchrony of neural discharge (signal morphology) and the non-Gaussian component (frequency coupling) is probably heterogeneous. The squared bicoherence values of normal EEG rhythms (Figures 6(a), 6(b), 6(c), 6(d), and 6(e)) show wide ranging peaks in

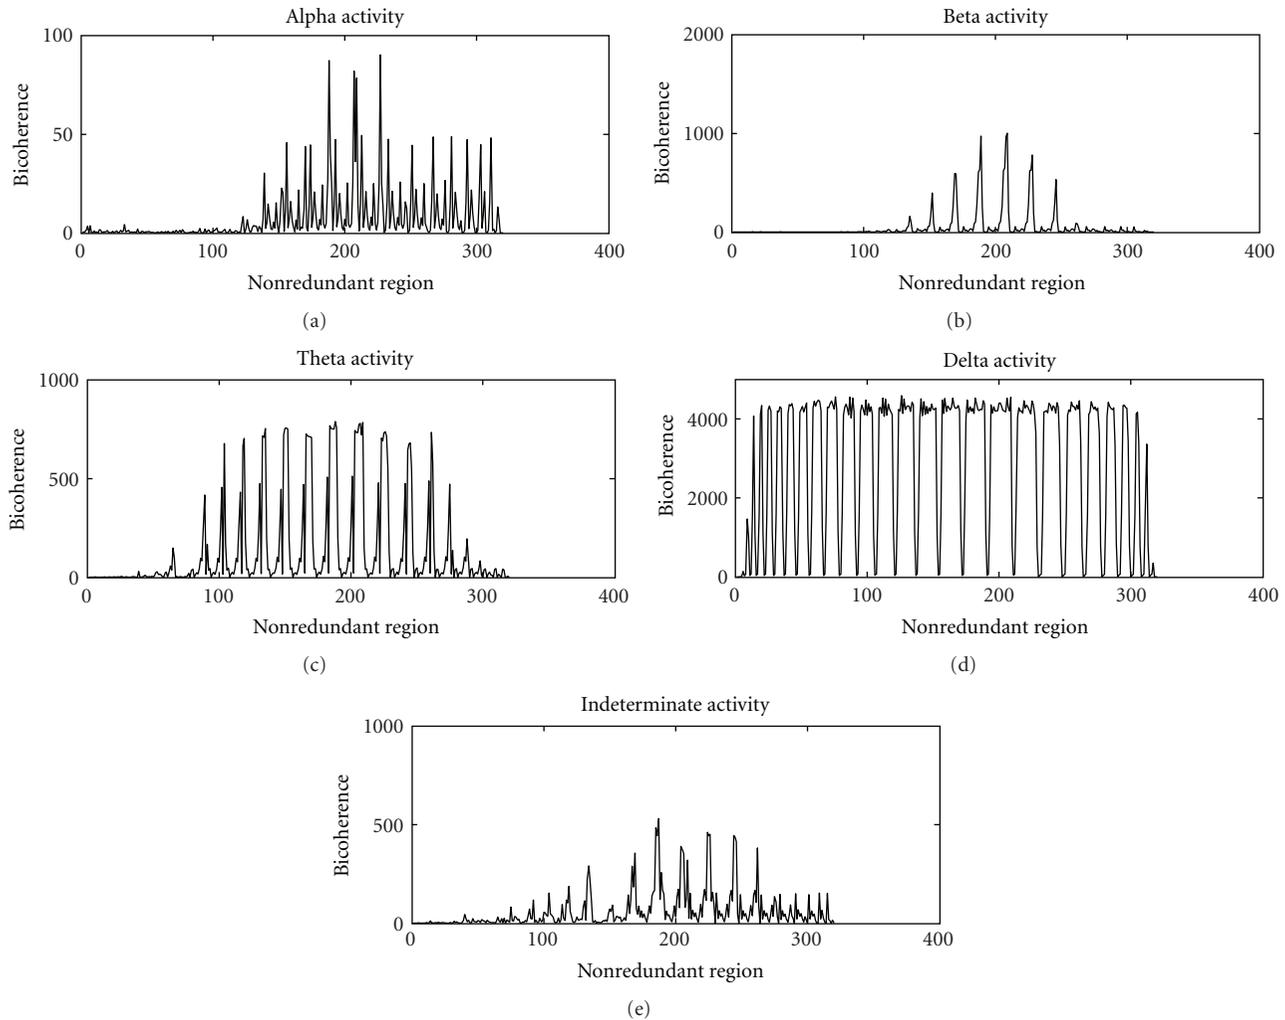


FIGURE 6: The bicoherence plots of (a) alpha EEG activity, (b) beta EEG activity, (c) theta EEG activity, (d) delta EEG activity, and (e) indeterminate EEG activity.

the nonredundant region. The coupling patterns, therefore, may be different for different EEG activity.

5. Conclusion

In our detection of nonlinearity or absence of linear stochastic mechanism, we have shown that the higher-order spectra can reliably distinguish chaotic signals and EEG rhythms from the filtered noises. The filtered noises in the lower passbands show linear stochastic properties in the distribution of their bicoherence values. The results of Hinich statistics and bicoherence estimates indicate that EEG rhythms have similar properties as those of the chaotic time series. The EEG signals are unequivocally non-Gaussian and nonlinear in character. In addition, the bicoherence patterns in the nonredundant regions of the EEG time series are similar to chaotic time series, reflecting quadratic phase coupling. The bispectrum preserves nonminimum phase information of a signal and outputs zero spectrum for linear mechanism. Therefore, the bicoherence statistics of nonredundant region of the

spectra (Hinich statistics) may be suitable for detecting hidden structure in signals [25].

References

- [1] T. H. Bullock, J. Z. Achimowicz, R. B. Duckrow, S. S. Spencer, and V. J. Iragui-Madoz, "Bicoherence of intracranial EEG in sleep, wakefulness and seizures," *Electroencephalography and Clinical Neurophysiology*, vol. 103, no. 6, pp. 661–678, 1997.
- [2] J. P. Pijn, J. Van Neerven, A. Noest, and F. H. Lopes Da Silva, "Chaos or noise in EEG signals; dependence on state and brain site," *Electroencephalography and Clinical Neurophysiology*, vol. 79, no. 5, pp. 371–381, 1991.
- [3] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. Doynne Farmer, "Testing for nonlinearity in time series: the method of surrogate data," *Physica D*, vol. 58, no. 1–4, pp. 77–94, 1992.
- [4] P. E. Rapp, I. D. Zimmerman, A. M. Albano, G. C. de Guzman, N. N. Greenbaum, and T. R. Bashore, "Experimental studies of chaotic neural behavior: cellular activity and electroencephalographic signals," in *Nonlinear Oscillations in Biology and*

- Chemistry*, H. G. Othmer, Ed., Springer, Berlin, Germany, 1985.
- [5] A. M. Albano, A. I. Mees, G.C. de Guzman, and P. E. Rapp, "Data requirements for reliable estimation of correlation dimensions," in *Chaos in Biological Systems*, H. Degn, A. V. Holden, and L. F. Olsen, Eds., Plenum, New York, NY, USA, 1987.
 - [6] T. M. McKenna, T. A. McMullen, and M. F. Shlesinger, "The brain as a dynamic physical system," *Neuroscience*, vol. 60, no. 3, pp. 587–605, 1994.
 - [7] E. Basar, *Chaos in Brain Function*, Springer, Berlin, Germany, 1990.
 - [8] I. Dvorak and A. V. Holden, *Mathematical Approaches to Brain Functioning Diagnostics*, Manchester University Press, 1991.
 - [9] B. H. Jansen, "Quantitative analysis of electroencephalograms: Is there chaos in the future?" *International Journal of Bio-Medical Computing*, vol. 27, no. 2, pp. 95–123, 1991.
 - [10] B. H. Jansen, "Is it and so what? A critical review of EEG chaos," in *Measuring Chaos in the Human Brain*, D. W. Duke and W. S. Pritchard, Eds., World Scientific, Singapore, 1991.
 - [11] P. E. Rapp, "Chaos in the neurosciences: cautionary tales from the frontier," *Biologist*, vol. 40, no. 2, pp. 89–94, 1993.
 - [12] P. E. Rapp, in *Chaos Theory in Psychology and Life Sciences*, A. Combs, Ed., Lawrence Erlbaum Associates, 1995.
 - [13] M. Palus, *Nonlinearity in normal human EEG: Cycles and randomness, not chaos*, Santa Fe Institute, Santa Fe, NM, USA, 1994, Santa Fe Institute Publication Number 94-10-054.
 - [14] W. S. Pritchard, D. W. Duke, and K. K. Kriebel, "Dimensional analysis of resting human EEG II: surrogate-data testing indicates nonlinearity but not low-dimensional chaos," *Psychophysiology*, vol. 32, no. 5, pp. 486–491, 1995.
 - [15] D. Prichard and J. Theiler, "Generating surrogate data for time series with several simultaneously measured variables," *Physical Review Letters*, vol. 73, no. 7, pp. 951–954, 1994.
 - [16] W. S. Pritchard and D. W. Duke, "Measuring chaos in the brain: a tutorial review of nonlinear dynamical EEG analysis," *International Journal of Neuroscience*, vol. 67, no. 1–4, pp. 31–80, 1992.
 - [17] N. Pradhan, P. K. Sadasivan, S. Chatterji, and D. Narayana Dutt, "Patterns of attractor dimensions of sleep EEG," *Computers in Biology and Medicine*, vol. 25, no. 5, pp. 455–462, 1995.
 - [18] P. E. Rapp, A. M. Albano, T. I. Schmah, and L. A. Farwell, "Filtered noise can mimic low-dimensional chaotic attractors," *Physical Review E*, vol. 47, no. 4, pp. 2289–2297, 1993.
 - [19] L. Glass, D. T. Kaplan, and J. E. Lewis, "Tests for deterministic dynamics in real and model neural networks," in *Nonlinear Dynamical Analysis of the EEG*, B. H. Jansen and M. E. Brandt, Eds., World Scientific, Singapore, 1992.
 - [20] S. J. Schiff, T. Sauer, and T. Chang, "Discriminating deterministic versus stochastic dynamics in neuronal activity," *Integrative Physiological and Behavioral Science*, vol. 29, no. 3, pp. 246–261, 1994.
 - [21] N. Pradhan and D. N. Dutt, "A nonlinear perspective in understanding the neurodynamics of EEG," *Computers in Biology and Medicine*, vol. 23, no. 6, pp. 425–442, 1993.
 - [22] Z. J. Kowalik and T. Elbert, "A practical method for the measurements of the chaoticity of electric and magnetic brain activity," *International Journal of Bifurcation and Chaos*, vol. 5, no. 2, pp. 475–490, 1995.
 - [23] B. P. Van Milligen, E. Sáncietz, T. Estrada et al., "Wavelet bicoherence: a new turbulence analysis tool," *Physics of Plasmas*, vol. 2, no. 8, pp. 3017–3032, 1995.
 - [24] A. Swami, J. Mendel, and C. Nikias, "Higher-Order Spectral analysis toolbox For use with Matlab (Mathworks 1995)".
 - [25] N. Pradhan, in *Nonlinear Dynamics and Brain Functioning*, N. Pradhan, P. E. Rapp, and R. Sreenivasan, Eds., Novascience Publishers, New York, NY, USA, 2000.
 - [26] S. Shahid and J. Walker, "Cepstrum of bispectrum—a new approach to blind system reconstruction," *Signal Processing*, vol. 88, no. 1, pp. 19–32, 2008.
 - [27] T. H. Bullock, J. Z. Achimowicz, R. B. Duckrow, S. S. Spencer, and V. J. Iragui-Madoz, "Bicoherence of intracranial EEG in sleep, wakefulness and seizures," *Electroencephalography and Clinical Neurophysiology*, vol. 103, no. 6, pp. 661–678, 1997.
 - [28] "Bispectral Index, BIS, BISx, and Aspect are trademarks or registered trademarks of Aspect Medical Systems".
 - [29] A. Saikia and S. M. Hazarika, "Bispectrum analysis of EEG during observation and imagination of hand movement," in *Proceedings of the IEEE Students' Technology Symposium (Tech-Sym '11)*, pp. 128–133, January 2011.
 - [30] O. Tirel, E. Wodey, R. Harris, J. Y. Bansard, C. Ecoffey, and L. Senhadji, "The impact of age on bispectral index values and EEG bispectrum during anaesthesia with desflurane and halothane in children," *British Journal of Anaesthesia*, vol. 96, no. 4, pp. 480–485, 2006.
 - [31] A. Gallego, J. F. Gil, J. M. Vico, E. Castro, J. E. Ruzzante, and R. Piotrkowski, "Wavelet transform and bispectrum applied to acoustic emission signals from adherence scratch-tests on corroded galvanized coatings," *Advanced Materials Research*, vol. 13-14, pp. 83–88, 2006.
 - [32] E. P. Tsolaki, "Testing nonstationary time series for Gaussianity and linearity using the evolutionary bispectrum: an application to internet traffic data," *Signal Processing*, vol. 88, no. 6, pp. 1355–1367, 2008.
 - [33] J. Ramírez, J. M. Górriz, and J. C. Segura, "Statistical voice activity detection based on integrated bispectrum likelihood ratio tests for robust speech recognition," *Journal of the Acoustical Society of America*, vol. 121, no. 5, pp. 2946–2958, 2007.

Research Article

Simulation of Spread and Control of Lesions in Brain

Krishna Mohan Thamattoor Raman

CSIR Centre for Mathematical Modelling and Computer Simulation (C-MMACS), Bangalore 560017, India

Correspondence should be addressed to Krishna Mohan Thamattoor Raman, kmohan@cmmacs.ernet.in

Received 19 September 2011; Accepted 25 October 2011

Academic Editor: Vikas Rai

Copyright © 2012 Krishna Mohan Thamattoor Raman. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

A simulation model for the spread and control of lesions in the brain is constructed using a planar network (graph) representation for the central nervous system (CNS). The model is inspired by the lesion structures observed in the case of multiple sclerosis (MS), a chronic disease of the CNS. The initial lesion site is at the center of a unit square and spreads outwards based on the success rate in damaging edges (axons) of the network. The damaged edges send out alarm signals which, at appropriate intensity levels, generate programmed cell death. Depending on the extent and timing of the programmed cell death, the lesion may get controlled or aggravated akin to the control of wild fires by burning of peripheral vegetation. The parameter phase space of the model shows smooth transition from uncontrolled situation to controlled situation. The simulations show that the model is capable of generating a wide variety of lesion growth and arrest scenarios.

1. Introduction

MS affects about one million people worldwide and causes physical and cognitive disability. There are three types of MS, relapsing-remitting, secondary progressive, and primary progressive, that differ in the dynamical patterns of disease progression. There are as yet no known cures for MS. Patients with relapsing MS are currently treated with drugs that exert immunomodulatory effects and slow the progression of the disease; there are no effective treatment options for the progressive forms of MS [1, 2].

MS is postulated to be a cell-mediated autoimmune disease directed against myelin components of the CNS. Myelin is an electrically insulating phospholipid layer that surrounds the axons of many neurons. The disease is characterized by both inflammatory immune responses and neurodegeneration. The prevailing hypothesis on MS pathogenesis is that autoreactive T-lymphocytes, a cell type in the immune system, orchestrate a complex cascade of events that cause blood-brain barrier disruption and invasion of immunologically aggressive cells into the CNS. However, the exact causes of MS still remain unknown [3, 4]. The long-term goals of this research are to develop disease models that can be used to evaluate therapeutic strategies for this disease and, in this report, the specific focus is on evaluating a network

model for MS lesion dynamics. Literature survey indicates that network approaches have not been studied extensively for disease modeling in MS.

1.1. Previous Work. Conventional models for autoimmunity are premised on the occurrence of defects in the immune system that cause it to turn against the host tissue. A defect-free immune system, in this world view, purportedly only attacks pathogens, the external agents that cause illness or disease [5–7]. However, an alternative viewpoint has been advocated where auto-immunity is seen as the usual immune response, but directed against those components of the body which, in normal conditions, are inaccessible to the immune system [8–14]. For example, in the *danger model*, developed by Matzinger [10, 11], it is posited that stressed and injured tissues can mediate immune responses through the generation of appropriate “danger” signals. This is as opposed to the activation through recognition of external pathogenic cell types from host tissue in the conventional models. The concept of *comprehensive immunity*, developed by Nevo et al. [12, 13] complements this alternate perspective; experimental results supporting their idea have also been reported [14]. The present network model is inspired by the alternative viewpoint.

The key elements of the model consist of a pathological process that causes cellular damage and programmed cell death (apoptosis) initiated through an intercellular signaling component. The programmed cell death deprives the pathological process of healthy tissue which is necessary for its propagation in space and time. In this, it resembles the action of firemen who burn peripheral vegetation to contain forest fires. Inter-cellular signaling is a key feature of the model that allow pathologically damaged cells to propagate alarm signals and initiate programmed cell death.

2. Model

An undirected, fixed radius random graph $G(n, r)$, with n nodes (vertices) and radius of connectivity, r , is constructed to represent the CNS in this 2D network model. Fixed radius implies that nodes are connected only if they are within a distance of r . Biologically, the nodes of the graph can be viewed as representing cell bodies or functional units and the edges (bonds) of the graph can be viewed as axons or the interconnections between functional units.

Let d_i be the degree of the i th node, that is, the number of edges attached to it. The health status of each edge, at time t , is indicated by its "weight," $w(j, t)$ ($j = 1, \dots, d_i$), an integer number ranging from $0 \dots w_{\max}$. Edges with weight w_{\max} are fully functional or healthy units (as at the beginning of simulation), and those with weight zero, are dead. Extending the same logic, the amplitude of the signal propagated along j th edge is taken to be equal to $w(j, t)$.

In the pathological process, the edges are damaged by lowering their weight by a single unit. However, in the programmed cell death process, edge weights are directly reduced to zero. In the regeneration process, edge weights are raised by a unit.

The pathological and regeneration processes are driven by probabilistic events wherein each edge in the affected region, in each time unit, has a certain probability $\mathbf{p}_i^d(\mathbf{p}_i^r)$ of getting damaged (regenerated). In the general case, $\mathbf{p}_i^d(\mathbf{p}_i^r)$ is a column vector of length w_{\max} containing the transition probabilities from one state of health to another. Probability of programmed cell death, p_p , is independent of the health status of the edge.

The functional or health status of the i th node is the sum over its edge weights, $s_i(t) = \sum_{j=1}^{d_i} w(j, t)$. The maximum possible value of s_i is denoted by S_i , which is realized when each $w(j, t) = w_{\max}$.

A node damaged by the pathological process generates an alarm signal when the ratio of its health status to the fully healthy state ($s_i(t)/S_i$) falls below a threshold, τ_{al} . The signals received at the i th node are summed and propagated further when the summed signal strength reaches s_i .

Programmed cell death is initiated at all the nodes where the propagated signals reach a threshold τ_{bf} . The accumulated alarm signals in the region of programmed cell death, a circular region around the activated node of radius proportional to a parameter C_{bf} , get reset to zero. No additional signals are generated at these nodes to the alarm signals generated in the pathologic process.

The spread of the pathologic process is driven by the success rate in causing cellular damage. The fraction of edges ($R_I(t)$) damaged in a particular time step, among the total number of healthy edges visited, is the rate of damage due to the pathologic process. The rates of damage due to the pathologic and the programmed cell death are computed in terms of the initial lesion size so that the final results are invariant with respect to the initial lesion size. Thus, the radius of the region affected by the pathologic process increased or decreased according to the formula, $\alpha \times R_I(t) \times ROI_{t=0}$, where $ROI_{t=0}$ is the radius of the region at the center where the initial lesion is seeded. In a similar fashion, the region of programmed cell death was computed as $C_{bf} \times ROI_{t=0}$.

3. Simulation

In the simulations reported here, a two-state model with $w_{\max} = 1$ has been employed, that is, there are no intermediate states of health, and the edges are either alive or dead. Additionally, the regeneration probability, \mathbf{p}_i^r , was set to zero in order to focus exclusively on the effects of the interplay between the pathological and programmed cell death processes on lesion structure and dynamics. A few preliminary results using such a configuration was reported earlier [15].

We have set $n = 400$ and chosen a uniform random distribution of points in the unit square $[0, 1] \times [0, 1]$. The radius of connectivity was set to $r = 0.2$. All the results were also confirmed on a network of $n = 4000$, with $r = 0.06$. Average degree strengths of the order of 10 are obtained in these configurations; degree distribution is Gaussian. The pathological process was initiated at $t = 0$ in a region with $ROI_{t=0} = 0.05$ around the center at $(0.5, 0.5)$; for $n = 4000$, $ROI_{t=0} = 0.015$.

A uniform probability of pathologic damage $p_d = 0.33$ was used, with $\alpha = 0.12$. We varied τ_{al} , τ_{bf} , and C_{bf} to identify the conditions under which the pathological process could be controlled by the programmed cell death. Larger values of τ_{bf} indicate reduced sensitivity to the alarm signals whereas a larger value of C_{bf} indicates that a larger area near the alerted node is subjected to programmed cell death. In the case of τ_{al} , larger values indicate quicker firing of alarm signals.

4. Results

Figure 1 shows the time series of damages caused to the system by both the pathological process and the programmed cell death process. The first column of panels in the figure shows the time course of instantaneous damages to the system. The middle column of panels shows the time course of the cumulative damages to the system. The last column of panels show the final state of the network at the end of the simulations.

There are three typical scenarios which are illustrated in Figure 1, in the three rows from top to bottom. Figures 1(a)–1(c) show a scenario where the programmed cell death is not of sufficient strength to significantly affect the pathologic process. Note that the instantaneous damages from

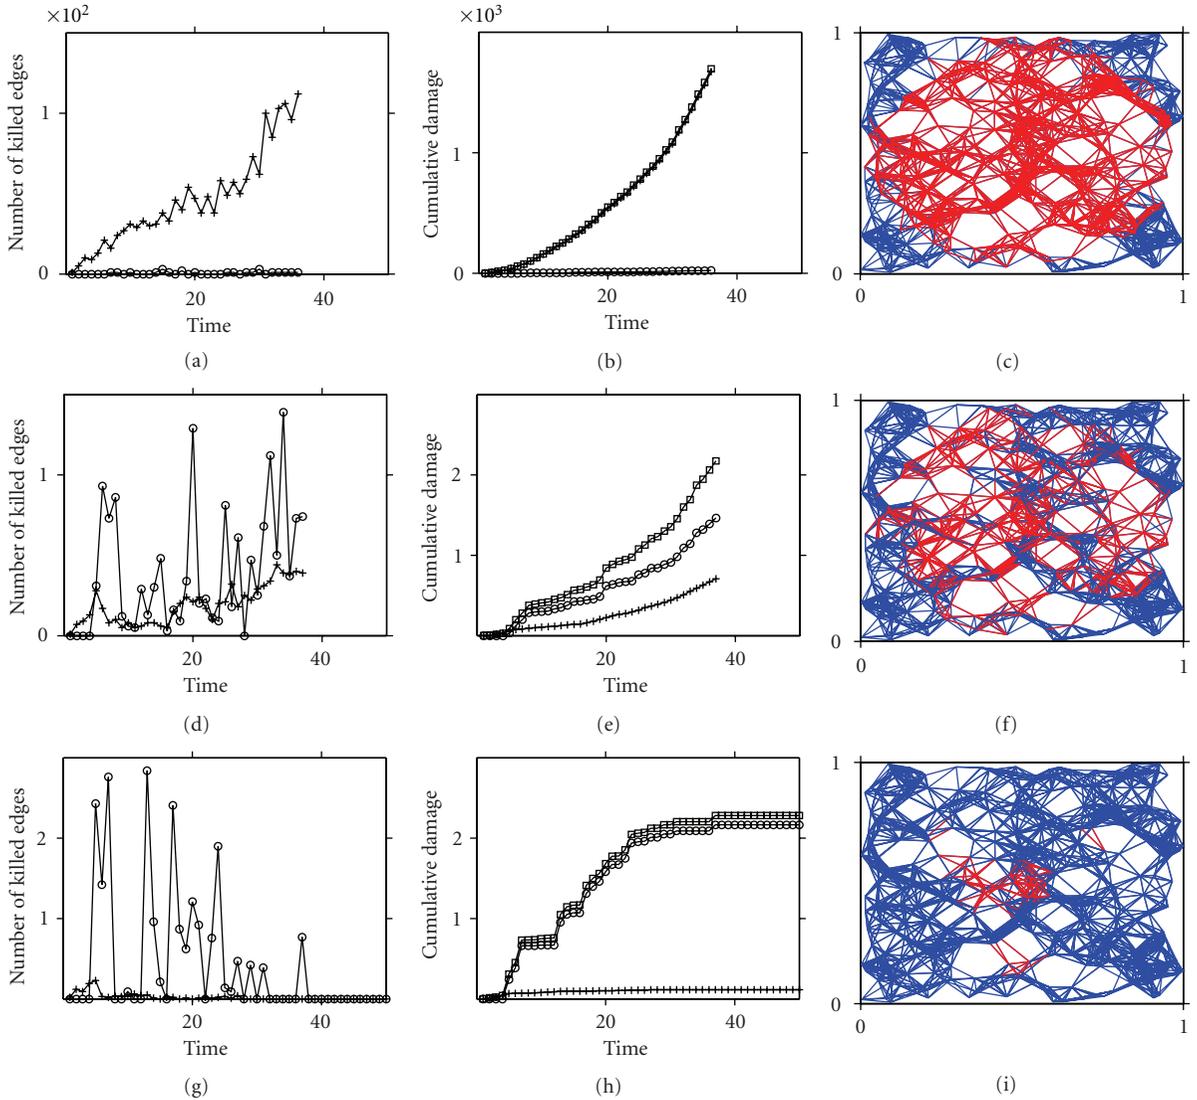


FIGURE 1: Time course of damages to the system by the pathological and programmed cell death processes. The first column of panels shows, separately, the instantaneous damages due to both the processes; dotted lines with asterisks indicate the damages due to the pathological process, and bold lines with filled circles indicate damages due to programmed cell death. The second column of panels shows the cumulative damages with time. Again, the damages effected through both the processes have been separately shown (same symbols as earlier), as also the sum total damages to the system (square symbol). The last column of panels (color online) show the state of the network at the end of the simulations; the dotted lines indicate healthy edges (axons), full (blue) lines indicate edges damaged due to programmed cell death, and dark (red) lines indicate edges damaged due to the pathological process. For all the panels, $\tau_{bf} = 0.5$ and $\tau_{al} = 0.7$, while the C_{bf} values, for each row, top to bottom, are 0.2, 0.8, and 1.5, respectively.

programmed cell death are hardly ever above zero. Also, it is seen from Figure 1(b) that the contribution of programmed cell death to the sum total of damages is insignificant. This situation occurs with a suitable combination of low τ_{al} , high τ_{bf} , and low C_{bf} values. Figures 1(d)–1(f) show a slightly more complex situation. In this case, programmed cell death is clearly the dominant effect. The instantaneous damages caused by both the processes are consistently nonzero (Figure 1(d)) and the cumulative damages (Figure 1(e)) continue to grow. The total damage, thus, continues to spread. In Figures 1(g)–1(i), the pathological process has been well controlled. The instantaneous damages have fallen to zero in

Figure 1(g), and the cumulative damages (Figure 1(h)) have leveled off. The final state of the network (Figure 1(i)) shows that the damage is also minimal in terms of the fraction of edges damaged.

As seen from Figures 1(a), 1(d), and 1(g), the time series is stochastic. There are essentially two sources of randomness in the model. Firstly, the pathological process is simulated by a binomial process wherein each edge visit will lead to successful damage if the generated random number falls below the value in \mathbf{p}_i^d for that edge. Secondly, the random network itself is generated by the random distribution of the n points in the plane. The complete picture of the transition

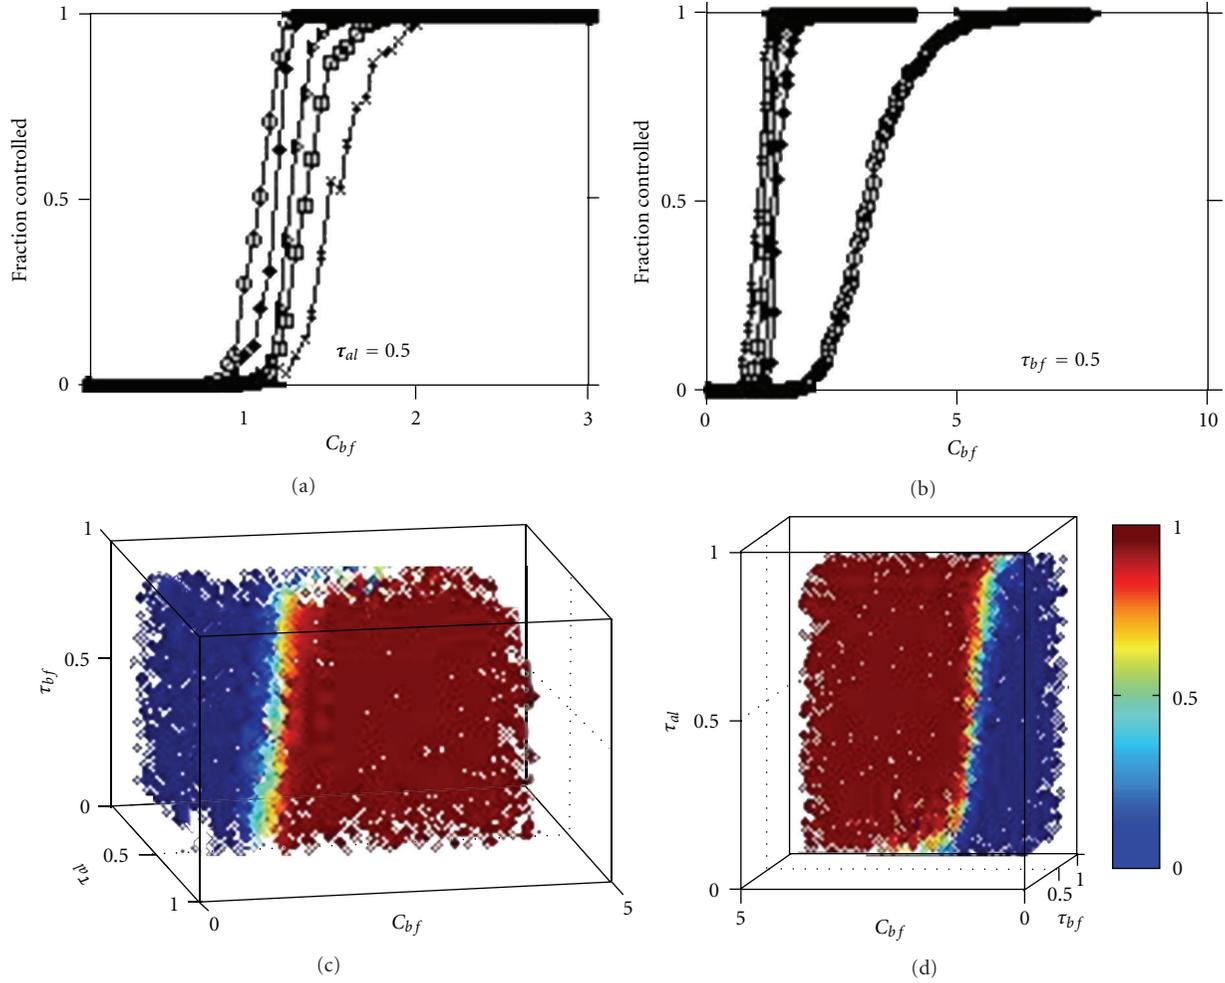


FIGURE 2: Phase portraits of the dynamics. The y -axes in Figures 2(a)-2(b) and the gray (color online) scale in Figures 2(c)-2(d) indicate the fraction of the total number of simulations where the pathological process was successfully arrested. In Figure 2a, τ_{al} has been held fixed at the shown value while the different curves are, from left to right, for $\tau_{bf} = 0.1, 0.3, 0.5, 0.7$, and 0.9 . In Figure 2(b), the different curves, from right to left, are for $\tau_{al} = 0.9, 0.7, 0.5, 0.3$, and 0.1 ; $\tau_{bf} = 0.5$ is held fixed. The S-curves that result with increasing C_{bf} values indicate the smooth nature of the transition in dynamical behavior in the parameter space.

from uncontrolled growth of the pathological process to the situation where the pathological process has been well arrested is seen in the parameter phase space graphs shown in Figure 2, where an averaging has been effected over the two sources of randomness. The phase space diagrams are the results of averaging over ten different networks, with the dynamics averaged over a thousand iterations.

From Figure 2, we see that the transition from uncontrolled pathological process to arrested pathological process is smooth as C_{bf} is varied from low to high values. In Figure 2(a), τ_{al} has been held fixed and the different curves, from left to right, are for different τ_{bf} values, from 0.1 to 0.9 in steps of 0.2 . In Figure 2(b), τ_{bf} has been held fixed and the different curves are, from right to left, for τ_{al} values ranging from 0.1 to 0.9 , in steps of 0.2 . We shall denote by $C_{bf}^{cr}(\tau_{al}, \tau_{bf})$ the critical value of C_{bf} at which these S-curves attain a value of 1 , that is, all instances of simulations result in the growth of pathological process being arrested. The

combined picture in the three parameter space is presented in Figures 2(c)-2(d), from two different perspectives. The three different scenarios presented, from top to bottom, in Figure 1 indicate, respectively, the three different parts, from left to right, of a typical S-curve of Figures 2(a)-2(b).

As seen from Figure 2, pathological process is always controlled if $C_{bf} > C_{bf}^{cr}$. Nevertheless, the sum total damage to the system is not the same for all values of $C_{bf} > C_{bf}^{cr}$; in fact, the damage is greater, the larger the value of C_{bf} . Clearly, it is desirable to effect control of the pathological process with the least sum total damage to the system. With this in mind, average fractional damages at different C_{bf} values have been plotted in Figure 3. Three different curves for three different τ_{bf} values are shown in this figure; similar graphs can be constructed for different τ_{al} values as well (not shown here; see [16]). These averages have been taken at $t = 20$ in each case. For $C_{bf} < C_{bf}^{cr}$ values, the damages due to the pathological process as well as sum total damage

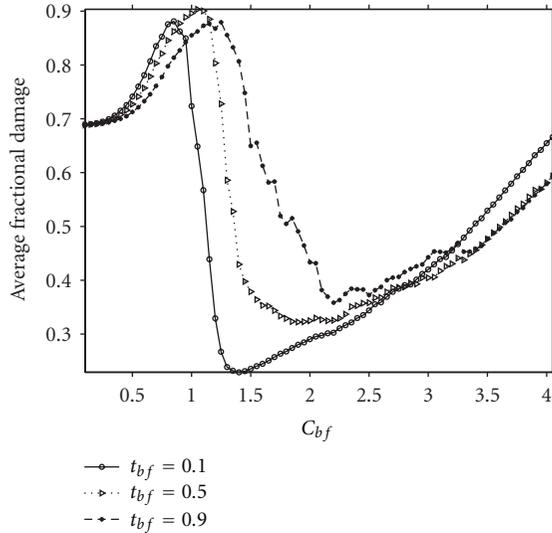


FIGURE 3: Optimality in sum total damage to the system while effecting arrest of pathological process is shown here. With increasing C_{bf} values the minimum in sum total damage to the system occurs at about the same value as the critical value at which the fraction of simulations in which the pathological process is arrested attains unity (cf. Figure 2(a)). This minimum occurs at higher values of C_{bf} with larger τ_{bf} , and the minimum value also shifts upward. A similar situation occurs with different τ_{al} values (not shown here). Note that the y -axis values have been taken at $t = 20$.

are still growing and have not become stationary at $t = 20$; for $C_{bf} > C_{bf}^{cr}$ values, the averages have become stationary. Nevertheless, these curves indicate that the least sum total damage to the system, with pathological process arrested, is obtained at $C_{bf} = C_{bf}^{cr}$. For $C_{bf} > C_{bf}^{cr}$, the programmed cell death is clearly effecting more damage than is necessary to arrest the pathological process. Since C_{bf}^{cr} depends on τ_{bf} and τ_{al} , it is not surprising to see that (cf. Figures 2(a) and (3)) lesser damage results when τ_{bf} is small.

From the above, it is clear (cf. Figures 1(d)-1(e)) that arrest of the pathological process does not necessarily occur if the damage due to the programmed cell death process is greater than the pathological process. What is necessary [16] is that the programmed cell death process be able to encircle the region affected by the pathological process, and, furthermore, be able to create an envelope region of sufficient thickness to offset its likely growth factor, $\alpha \times R_I(t) \times ROI_{t=0}$. This is achieved in all instances of simulation when $C_{bf} > C_{bf}^{cr}$. Currently, mathematical analysis of this feature is being carried out to establish the relationship of C_{bf}^{cr} with τ_{al} and τ_{bf} , and the results will be reported soon.

5. Conclusions

A physically motivated 2D network model was developed for the CNS and employed to study the process of lesion formation and spread in MS. Intercellular signalling of distress by the damaged cells is a key feature of the model which leads to programmed cell death getting activated in

an attempt to arrest the lesion progress. The model demonstrates that the spread of the pathological process can be arrested by programmed cell death when the geometry of the damage inflicted by the latter leads to an envelope, of sufficient thickness, being created encircling the area of pathological process. Such an envelope of dead cells deprives the pathological process of healthy cells which can sustain its growth. The model shows a smooth transition, as parameters are varied, from the situations of run-away pathological process, through aggravated damage to the system caused by unsuccessful firing of programmed cell death, to the creation of successful envelope around the pathological process.

The model complements the alternate viewpoint on autoimmunity which posits that cells and tissues signal distress and activate the immune system. Such a viewpoint circumvents the need for the immune system to store information about likely pathogens and, also, makes it capable of acting in instances of cellular damage resulting from nonpathogenic causes. Further study of the model along with identification of the possible biological constituents should enable comparisons with experiments and a more detailed exposition.

Acknowledgments

This work was carried out during a period spent at the Department of Pharmaceutical Sciences, State University of New York at Buffalo (UB), USA. Encouragement, discussions, and funding from Dr. Murali Ramanathan are acknowledged. Use of the computational facilities at the Center for Computational Research, UB, facilitated the work.

References

- [1] O. Kantarci and D. Wingerchuk, "Epidemiology and natural history of multiple sclerosis: new insights," *Current Opinion in Neurology*, vol. 19, no. 3, pp. 248–254, 2006.
- [2] T. J. Murray, "Diagnosis and treatment of multiple sclerosis," *British Medical Journal*, vol. 332, no. 7540, pp. 525–527, 2006.
- [3] M. Sospedra and R. Martin, "Immunology of multiplesclerosis," *Annual Review of Immunology*, vol. 23, pp. 683–747, 2005.
- [4] V. H. Perry and D. C. Anthony, "Axon damage and repair in multiple sclerosis," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 354, no. 1390, pp. 1641–1647, 1999.
- [5] F. M. Burnet and F. Fenner, *The Production of Antibodies*, MacMillan, Melbourne, VIC, Australia, 2nd edition, 1949.
- [6] F. M. Burnet, Ed., *Immunology: Readings from Scientific American*, W. H. Freeman, San Francisco, Calif, USA, 1976.
- [7] E. Christ and A. I. Tauber, "Selfhood, immunity, and the biological imagination: the thought of Frank Macfarlane Burnet," *Biology and Philosophy*, vol. 15, pp. 509–533, 1999.
- [8] G. J. V. Nossal, *Antibodies and Immunity*, Basic Books, New York, NY, USA, 2nd edition, 1978.
- [9] O. A. Smirnova and N. V. Stepanova, "Mathematical model of autoimmunity," *Biofizika*, vol. 20, no. 6, pp. 1095–1098, 1975.
- [10] P. Matzinger, "Tolerance, danger, and the extended family," *Annual Review of Immunology*, vol. 12, pp. 991–1045, 1994.
- [11] P. Matzinger, "An innate sense of danger," *Seminars in Immunology*, vol. 10, no. 5, pp. 399–415, 1998.

- [12] U. Nevo, J. Kipnis, I. Golding et al., “Autoimmunity as a special case of immunity: removing threats from within,” *Trends in Molecular Medicine*, vol. 9, no. 3, pp. 88–93, 2003.
- [13] U. Nevo, I. Golding, A. U. Neumann, M. Schwartz, and S. Akselrod, “Autoimmunity as an immune defense against degenerative processes: a primary mathematical model illustrating the bright side of autoimmunity,” *Journal of Theoretical Biology*, vol. 227, no. 4, pp. 583–592, 2004.
- [14] G. Moalem, R. Leibowitz-Amit, E. Yoles, F. Mor, I. R. Cohen, and M. Schwartz, “Autoimmune T cells protect neurons from secondary degeneration after central nervous system axotomy,” *Nature Medicine*, vol. 5, no. 1, pp. 49–55, 1999.
- [15] T. R. Krishna Mohan, S. Sen, and M. Ramanathan, “A computational model for lesion dynamics in multiple sclerosis of the brain,” *International Journal of Modern Physics E*, vol. 17, no. 5, pp. 930–939, 2008.
- [16] T. R. Krishna Mohan, “Unpublished results”.

Research Article

Free Energy, Value, and Attractors

Karl Friston¹ and Ping Ao^{2,3}

¹The Wellcome Trust Centre for Neuroimaging, UCL, Institute of Neurology, 12 Queen Square, London WC1N 3BG, UK

²Shanghai Center for Systems Biomedicine, Key Laboratory of Systems Biomedicine of Ministry of Education, Shanghai Jiao Tong University, Shanghai 200240, China

³Departments of Mechanical Engineering and Physics, University of Washington, Seattle, WA 98195, USA

Correspondence should be addressed to Karl Friston, k.friston@fil.ion.ucl.ac.uk

Received 23 August 2011; Accepted 7 September 2011

Academic Editor: Vikas Rai

Copyright © 2012 K. Friston and P. Ao. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

It has been suggested recently that action and perception can be understood as minimising the free energy of sensory samples. This ensures that agents sample the environment to maximise the evidence for their model of the world, such that exchanges with the environment are predictable and adaptive. However, the free energy account does not invoke reward or cost-functions from reinforcement-learning and optimal control theory. We therefore ask whether reward is necessary to explain adaptive behaviour. The free energy formulation uses ideas from statistical physics to explain action in terms of minimising sensory surprise. Conversely, reinforcement-learning has its roots in behaviourism and engineering and assumes that agents optimise a policy to maximise future reward. This paper tries to connect the two formulations and concludes that optimal policies correspond to empirical priors on the trajectories of hidden environmental states, which compel agents to seek out the (valuable) states they expect to encounter.

1. Introduction

This paper is about the emergence of adaptive behaviour in agents or phenotypes immersed in an inconstant environment. We will compare and contrast two perspectives; one based upon a free energy principle [1] and the other on optimal control and reinforcement-learning [2–5]. The key difference between these perspectives rests on what an agent optimises. The free energy principle assumes that both the action and internal states of an agent minimise the surprise (the negative log-likelihood) of sensory states. This surprise does not have to be learned because it defines the agent. In brief, being a particular agent induces a probability density on the states it can occupy (e.g., a fish in water) and, implicitly, surprising states (e.g., a fish out of water). Conversely, in reinforcement-learning, agents try to optimise a policy that maximises expected reward. We ask how free energy and policies are related and how they specify adaptive behaviour. Our main conclusion is that policies can be cast as beliefs about the state-transitions that determine free energy. This has some important implications for understanding the quantities that the brain has to represent when responding adaptively to changes in the sensorium.

We have shown recently that adaptive behaviour can be prescribed by prior expectations about sensory inputs, which action tries to fulfill [6]. This is called *active inference* and can be implemented, in the context of supervised learning, by exposing agents to an environment that enforces desired motion through state-space [7]. These trajectories are learned and recapitulated in the absence of supervision. The resulting behaviour is robust to unexpected or random perturbations and can be used to solve benchmark problems in reinforcement-learning and optimal control: see [7] for a treatment of the mountain-car problem. Essentially, active inference replaces value-learning with perceptual learning that optimises empirical (acquired) priors in the agent's internal model of its world. These priors specify the free energy associated with sensory signals and guide action to ensure sensations conform to prior beliefs. In this paper, we consider the harder problem addressed by reinforcement-learning and other semisupervised schemes. These schemes try to account for adaptive behaviour, given only a function that labels states as attractive or costly. This means agents have to access distal attractors, under proximal constraints furnished by the environment and their repertoire of allowable actions. We will take a dynamical perspective on this

problem, which highlights the relationship between active inference and reinforcement-learning and the connection between empirical priors and policies.

This paper comprises five sections. The first considers adaptive behaviour in terms of equilibria and random attractors, which we attempt to link later to concepts in behavioural economics and optimal decision or game theory [8, 9]. This section considers autopoietic (self-creating) attractors to result from minimising the conditional entropy (average surprise) of an agent's states through action. However, agents can only infer hidden states of the environment given their sensory states, which means agents must minimise the surprise associated with sensations. The second section shows how agents can do this using an upper (free energy) bound on sensory surprise. This leads to a free energy formulation of well-known inference and learning schemes based on generative models of the world [10–13]. In brief, the imperatives established in the first section are satisfied when action and inference minimise free energy. However, the principle of minimising free energy also applies to the form of the generative model entailed by an agent (its formal priors). These encode prior beliefs about the transitions or motion of hidden states and ensuing attractors, which action tries to fulfil. These priors or policies are considered from a dynamical perspective in the remaining sections. Section three considers some universal policies, starting with the Helmholtz decomposition and introducing the notion of value, detailed balance, and divergence-free flow. The final two sections look at *fixed-point* and *itinerant* policies, respectively. Fixed-point policies attract trajectories to (low-cost) points in state-space. These policies are considered in reinforcement-learning and optimal control theory [2, 4, 14]. They are based on Lyapunov (value) functions that specify the policy. However, under the Helmholtz decomposition, value functions are an incomplete specification of policies. This speaks to more general forms of (itinerant) policies that rest on the autovivitation (self-destruction) of costly attractors and itinerant (wandering or searching) motion through state-space. We illustrate the basic ideas using the same mountain-car problem that we have used previously in the context of supervised learning [7].

The main conclusion of this paper is that it is sufficient to minimise the average surprise (conditional entropy) of an agent's states to explain adaptive behaviour. This can be achieved by policies or empirical priors (equations of motion) that guide action and induce random attractors in its state-space. These attract agents to (low-cost) invariant sets of states and lead to autopoietic and ergodic behaviour.

2. Ensemble Dynamics and Random Attractors

What do adaptive agents optimise? We address this question using an ensemble density formulation, which has close connections to models of evolutionary processes [15–17] and equilibria in game theory [18]. We also introduce a complementary perspective based on random dynamical systems [19]. The equilibrium approach rests on an ensemble density over the states of an agent. This can be regarded as the density of innumerable copies of the agent, each represented by

a point in phase or state-space. This density is essentially a probabilistic definition of the agent, in terms of the states it occupies. For a well-defined agent to exist its ensemble density must be *ergodic*; that is, an invariant probability measure [20]. In other words, the density cannot change over time; otherwise, the definition of an agent (in terms of the states it occupies) would change. A simple example here would be the temperature of an organism, whose ensemble density is confined to certain phase-boundaries. Transgressing these boundaries would change the agent into something else (usually a dead agent). The simple fact that an agent's ensemble density exists and is confined within phase-boundaries (i.e., is ergodic or invariant) has some fundamental implications, which we now consider more formally.

2.1. Set Up: States and Dependencies. If an agent and its environment have states, what does it mean for the states of an agent to be distinct from those of its environment? We will take this to mean that an agent has *internal* and *external* states that are conditionally independent and are therefore separated by a Markov blanket. The minimal (nontrivial) requirement for this blanket to exist is a partition of the states into two pairs of subsets, where one pair constitutes a Markov blanket for the other.

This straightforward consideration suggests a four-way partition of state-space $X \times S \times A \times M \subset \mathbb{R}$ associated with an agent $m \in M$. Here, external states $\tilde{x} \in X$ represent states of the agent's immediate environment, such as forces, temperature, and physiological states. The tilde notion denotes a generalised state, which includes temporal derivatives to arbitrarily high order, such that $\tilde{x} = [x, x', x'', \dots]^T$ comprises position, velocity, acceleration, jerk, and so on. The internal states $\tilde{\mu} \in M$ correspond to things like intracellular concentrations, neuronal activity, and so forth. We will see later that these are internal representations of external states. These states are separated from each other by a Markov blanket $S \times A$, comprising *sensory* states that mediate the influence of external states on internal states and *action*, which mediates the influence of internal states on external states. Sensory states $\tilde{s} \in S$, like photoreceptor activity, depend on external states, while action $a \in A$, like alpha motor neuron activity, depends on internal states. Figure 1 illustrates these conditional dependencies in terms of a graphical model, in which action and sensation form a Markov blanket separating external and internal states. In other words, external states are "hidden" from the agent's internal states. We will therefore refer to external states as *hidden* states.

The notion of a Markov blanket refers to a (statistical) boundary between the internal and hidden states of an agent. For simple (cellular) organisms, this could be associated with the cell surface, where sensory states correspond to the states of receptors and ion channels and action to various transporter and cell adhesion processes. For more complicated multicellular organisms (like us) the boundary of an agent is probably best thought of in terms of systems. For example, neuronal systems have clearly defined sensory states at their receptors and action is mediated by a discrete number of effectors. Here, the notion of a surface is probably less useful,

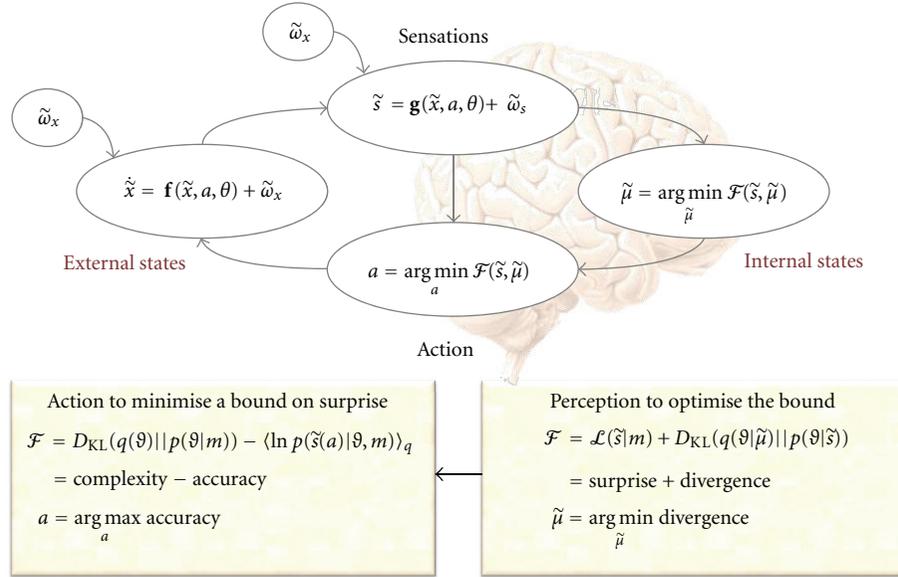


FIGURE 1: The free energy principle. The schematic shows the probabilistic dependencies (arrows) among the quantities that define free energy. These include the internal states of the brain $\tilde{\mu}(t)$ and quantities describing its exchange with the environment. These are the generalized sensory states $\tilde{z}(t) = [s, s', s'', \dots]^T$ and action $a(t)$. The environment is described by equations of motion, which specify the trajectory of its hidden states and a mapping to sensory states. The quantities $\vartheta \supset (\tilde{x}, \theta)$ causing sensory states comprise hidden states and parameters. The hidden parameters control the equations (\mathbf{f}, \mathbf{g}) and precision (inverse variance) of random fluctuations ($\omega_x(t), \omega_s(t)$) on hidden and sensory states. Internal brain states and action minimize free energy $\mathcal{F}(\tilde{z}, \tilde{\mu})$, which is a function of sensory states and a probabilistic representation $q(\vartheta | \tilde{\mu})$ of their causes. This representation is called the recognition density and is encoded by internal states that play the role of sufficient statistics. The free energy depends on two probability densities; the recognition density, $q(\vartheta | \tilde{\mu})$, and one that generates sensory samples and their causes, $p(\tilde{z}, \vartheta | m)$. The latter represents a probabilistic generative model (denoted by m), whose form is entailed by the agent. The lower panels provide alternative expressions for the free energy to show what its minimization entails. Action can only reduce free energy by increasing accuracy (i.e., selectively sampling sensory states that are predicted). Conversely, optimizing internal states makes the representation an approximate conditional density on the causes of sensory states. This enables action to avoid surprising sensory encounters. See main text for further details.

in the sense that the spatial deployment of sensory epithelia becomes a hidden state (and depends on action).

The external state-space we have in mind is high dimensional, covering the myriad of macroscopic states that constitute an embodied agent and its proximal environment. We assume that this system is open and that its states are confined to a low-dimensional manifold $\mathcal{O} \subset X$ that endow the agent with attributes. More precisely, the agent has *observables* (i.e., phenotypic traits or characteristics) that are given by real-valued functions, whose domain is the bounded set $\mathcal{O} \subset X$. This implies that there are states $\tilde{x} \notin \mathcal{O}$ an agent cannot occupy (e.g., very low temperatures). An observable is a property of the state that can be determined by some operator. A simple example of a bounded operator would be length, which must be greater than zero.

The existence of macroscopic states appeals to the fact that interactions among microscopic states generally lead to macroscopic order. There are many examples of this in the literature on complex systems and self-organisation. Key examples of macroscopic states are the order parameters used to describe phase-transitions [21]. The order parameter concept has been generalized to the slaving principle [22], under which the fast (stable) dynamics of rapidly dissipating patterns (modes or phase-functions) of microscopic states are determined by the slow (unstable) dynamics of a few macro-

scopic states (order parameters). These states can be regarded as the amplitudes of patterns that determine macroscopic behaviour. The enslaving of stable patterns by macroscopic states greatly reduces the degrees of freedom of the system and leads to the emergence of macroscopic order (e.g., pattern formation). A similar separation of temporal scales is seen in centre manifold theory [23]. See [24–26] for interesting examples and applications. We will assume that macroscopic states $\tilde{x} \in X$ are (unique phase) functions of the microscopic states that they enslave.

The emergence of macroscopic order (and its associated states) is easy to simulate. Figure 2 provides a simple example where sixteen (Lorenz) oscillators have been coupled to each other, so that each oscillator (with three microscopic states) sees all the other oscillators. In this example, the macroscopic states (c.f. order parameters) are just the average of each state over oscillators; this particular phase-function is known as a mean field; see [27] for a discussion of mean field treatments of neuronal dynamics. Here, the mean field enslaves the states of each oscillator so that the difference between each microscopic state and its average decays quickly; these differences are the stable patterns and decay to zero. This draws the microscopic states to a low- (three-) dimensional manifold, known as a synchronisation manifold [28]. Although the emergence of order is easy to simulate, it is also easy to

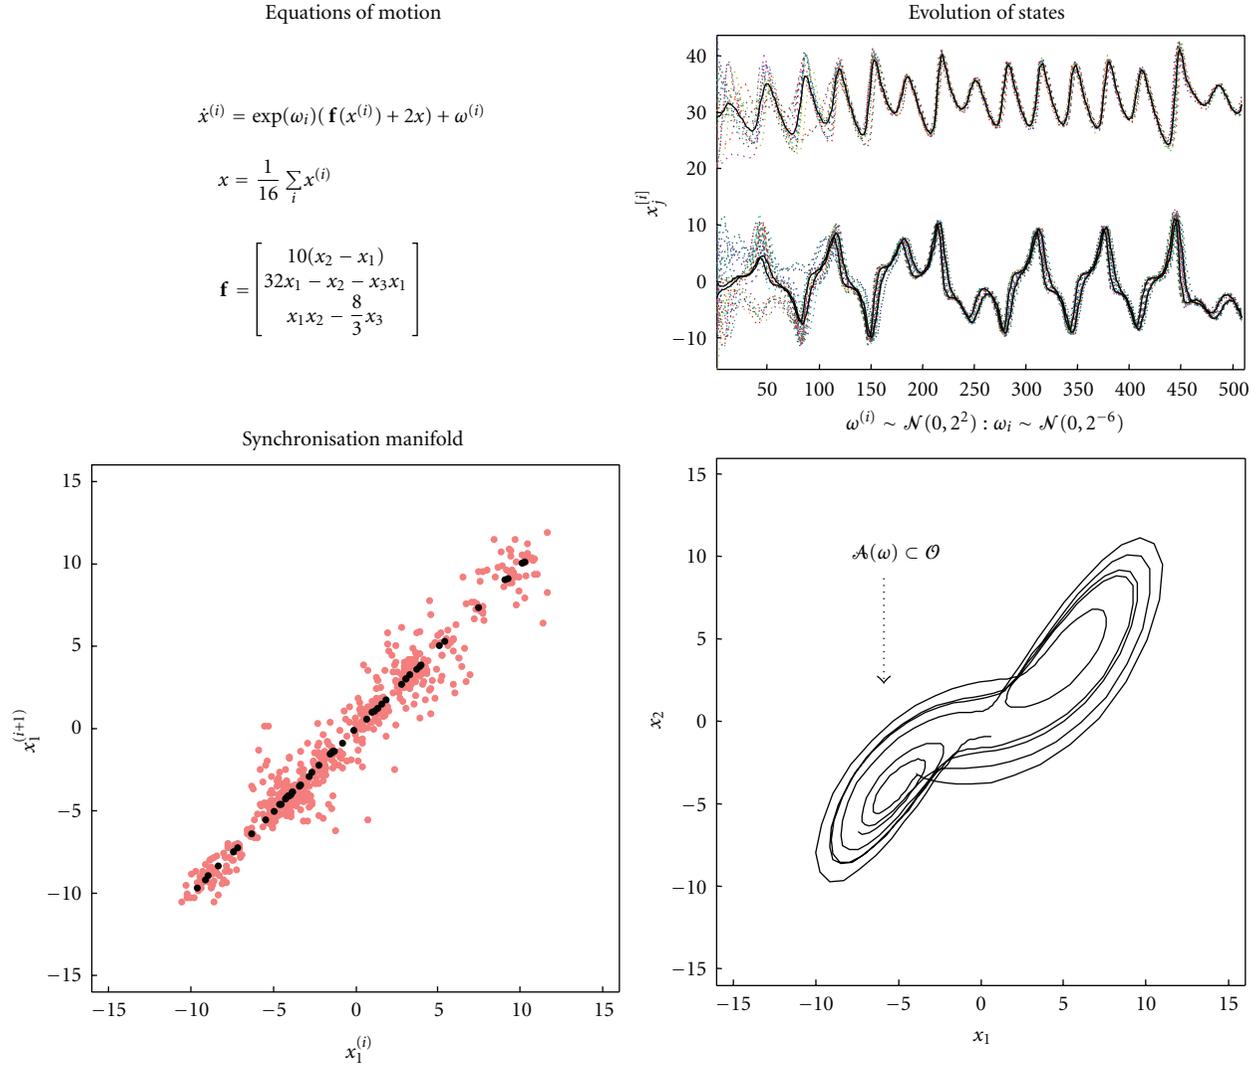


FIGURE 2: Self-organisation and the emergence of macroscopic behaviour. This figure shows a simple example of self-organisation using sixteen (Lorenz) oscillators that have been coupled to each other, so that each oscillator (with three microscopic states) sees the other oscillators. This is an example of a globally coupled map, where the dynamics of each oscillator conform to a classical Lorenz system. The equations of motion are provided in the left panel for each microstate, $x_j^{(i)} : i = 1, \dots, 16 : j = 1, 2, 3$, whose average constitutes a macrostate $x_j : j \in 1, 2, 3$. Each oscillator has its own random fluctuations $\omega^{(i)}(t) \in \mathbb{R}$ and speed $\exp(\omega_i) \in \mathbb{R}^+$. The upper right panel shows the evolution of the microstates (dotted lines) and the macrostates (solid lines) over 512 time steps of one 1/32 second. The lower right panel, shows the first two macrostates plotted against each other to show the implicit attractor that emerges from self-organisation. The lower left panel shows the implicit synchronisation manifold by plotting the first states from successive pairs of oscillators (pink) and their averages (black) against each other. This simulation used low levels of noise on the motion of the microstates $\omega^{(i)} \sim \mathcal{N}(0, 2^2)$ and the log-rate constants $\omega_i \sim \mathcal{N}(0, 2^{-6})$ that disperse the speeds of each oscillator. The initial states were randomised by sampling from a Gaussian distribution with a standard deviation of eight.

destroy. Figure 3 shows how macroscopic order collapses when the random fluctuations on the motion of states are increased. Here, there is no slaving because the system has moved from a coherent regime to an incoherent regime, where each oscillator pursues its own path. Order can also be destroyed by making the coherence trivial; this is known as oscillator death and occurs when each oscillator approaches a fixed-point in state-space (interestingly these fixed-points are unstable when the oscillators are uncoupled, see [24]). Oscillator death is illustrated in Figure 3 by increasing the random dispersion of speeds along each oscillators orbit

(trajectory). In these examples, macroscopic order collapses into incoherent or trivially coherent dynamics. We have deliberately chosen to illustrate these phenomena with a collection of similar oscillators (known technically as a globally coupled map; see also [29]), because the macroscopic dynamics recapitulate the dynamics of each oscillator in isolation. This means one could imagine that the microscopic states are themselves phase-functions of micromicroscopic states and so on *ad infinitum*. Heuristically, this speaks to the hierarchical and self-similar dynamics of complex self-organising systems [30, 31].

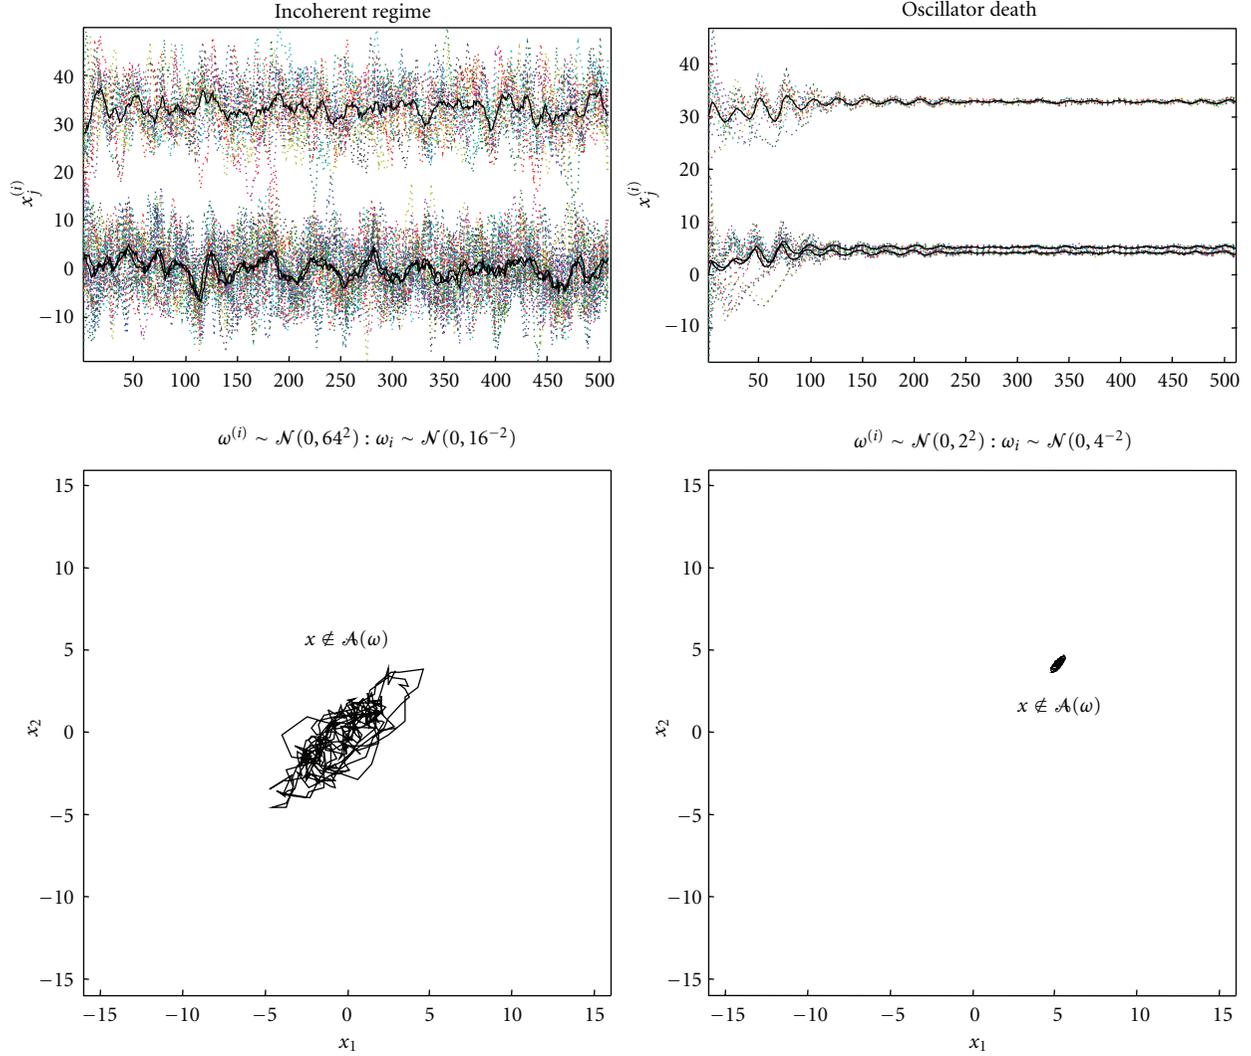


FIGURE 3: The loss of macroscopic order and oscillator death. This figure uses the same format and setup as in the previous figure but here shows the loss of macroscopic order through incoherence (left) and oscillator death (right). Incoherence was induced by increasing the random fluctuations on the motion of states to $\omega^{(i)} \sim \mathcal{N}(0, 2^{10})$. Oscillator death was induced by increasing the random dispersion of speeds along each oscillators orbit to $\omega_i \sim \mathcal{N}(0, 2^{-4})$, see [24]. The ensuing macroscopic states (lower panels) now no longer belong to the attracting set of the previous figure: $\mathcal{A}(\omega) \subset \mathcal{O}$.

In summary, the emergence of macroscopic order is not mysterious and arises from a natural separation of temporal scales that is disclosed by some transformation of variables. However, the ensuing order is delicate and easily destroyed. In what follows, we shall try to understand how self-organisation keeps the macroscopic states of an agent within a bounded set $\mathcal{O} \subset X$ for extended periods of time. To do this we will look more closely at their dynamics.

2.2. Dynamics and Ergodicity. Let the conditional dependencies among the (macroscopic) states $X \times S \times A \times M \subset \mathbb{R}$ in Figure 1 be described by the following coupled differential equations:

$$\begin{aligned} \dot{\tilde{x}} &= \mathbf{f}(\tilde{x}, a, \theta) + \tilde{\omega}_a, \\ \dot{\tilde{s}} &= \mathbf{g}(\tilde{x}, a, \theta) + \tilde{\omega}_s, \end{aligned} \quad (1)$$

where (as we will see later)

$$\begin{aligned} \dot{a} &= -\partial_a \mathcal{F}(\tilde{s}, \tilde{\mu}), \\ \dot{\tilde{\mu}} &= -\partial_{\tilde{\mu}} \mathcal{F}(\tilde{s}, \tilde{\mu}) + \mathcal{D}\tilde{\mu}. \end{aligned} \quad (2)$$

Here, \mathcal{D} is a derivative matrix operator with identity matrices along its first diagonal such that $\mathcal{D}\tilde{\mu} = [\mu', \mu'', \mu''', \dots]^T$. The first (stochastic differential) equation above describes the flow of hidden states in terms of a mapping $\mathbf{f} : X \times A \rightarrow X$ and some random fluctuations, $\tilde{\omega}_a \in \Omega_a$, while the second expresses sensory states in terms of a sensory mapping $\mathbf{g} : X \rightarrow S$ and noise, $\tilde{\omega}_s \in \Omega_s$. In this formulation, sensations are a noisy map of hidden states that evolve as a function of themselves and action, where exogenous influences from outside the proximal environment are absorbed into the random fluctuations. The quantities θ represent time-invariant

parameters of the equations of motion and sensory mapping. For simplicity, we will omit θ for the remainder of this section and return to them later. The second pair of equations describes action $a : M \times S \rightarrow A$ and internal states $\tilde{\mu} : M \times S \rightarrow M$ as a gradient descent on a functional (function of a function) of sensory and internal states: $\mathcal{F}(\tilde{s}, \tilde{\mu}) \in \mathbb{R}$. The purpose of this paper is to motivate the nature of this (free energy) functional and relate it to classical treatments of optimal behaviour.

As it stands, (1) is difficult to analyse because flow is a nonautonomous function of action. We can finesse this (without loss of generality) by expressing action as a function of the current state $u(\tilde{x}(t))$ plus a fluctuating part $\omega_u(t)$ using a Taylor expansion around the action expected in state $\tilde{x} \in X$

$$\begin{aligned}\tilde{\dot{x}} &= \mathbf{f}(\tilde{x}, a) + \tilde{\omega}_a \\ &= \mathbf{f}(\tilde{x}, u) + \tilde{\omega}_x \\ \tilde{\omega}_x &= \tilde{\omega}_a + \partial_u \mathbf{f} \cdot \omega_u + \dots \\ a(t) &= u(\tilde{x}) + \omega_u.\end{aligned}\tag{3}$$

Equation (3) reformulates the dynamics in terms of controlled flow $\mathbf{f}(\tilde{x}, u) := \mathbf{f} : X \rightarrow X$ and controlled fluctuations $\tilde{\omega}_x \in \Omega_x$. This formulation is autonomous in the sense that controlled flow depends only on the current state. Furthermore, it allows us to connect to the optimal control literature that usually assumes *control* $u(\tilde{x})$ is a function of, and only of, the current state. In our setup, control is the expected (average) action in a hidden state. In contrast, action $a : M \times S \rightarrow A$ depends on internal and sensory states and therefore depends upon hidden states and random fluctuations in the past. In what follows, we will refer to controlled flow as a *policy* in the sense that it describes motion through state-space or transitions among states, in the absence of random effects. The policy is also the expected flow because it is the flow under expected action.

With these variables in place we can now ask what can be deduced about the nature of action and control, given the existence of agents. Our starting point is that agents are ergodic [20, 32], in the sense that their ensemble density is invariant (conserved) over a suitably long time scale. This is just another way of saying that agents occupy a subset of states $\mathcal{O} \subset X$ for long periods of time. The implicit ergodic (invariant) density $p(\tilde{x} | m) := p(\tilde{x}, \infty | m)$ is the stationary solution to the Fokker-Planck equation (also known as the Kolmogorov forward equation; [33]) describing the dynamics of the ensemble density over hidden states

$$\begin{aligned}\dot{p}(\tilde{x}, t | m) &= \Lambda p := \nabla \cdot \Gamma \nabla p - \nabla \cdot (\mathbf{f} p) \\ \dot{p}(\tilde{x} | m) &= 0 \implies \\ p(\tilde{x} | m) &= \mathcal{E}(\Lambda).\end{aligned}\tag{4}$$

Here, $\Lambda(\mathbf{f}, \Gamma)$ is the Fokker-Planck operator and Γ is half the covariance (amplitude) of the controlled fluctuations (a.k.a. the diffusion tensor). Equation (4) assumes the fluctuations are temporally uncorrelated (Wiener) processes; however, because the fluctuations $\tilde{\omega}_x(t)$ are in generalised coordinates

of motion, the fluctuations on states *per se* can be smooth and analytic [34]. The Fokker-Planck equation exploits the fact that the ensemble (probability mass) is conserved. The first (diffusion) term of the Fokker-Planck operator reflects dispersion due to the fluctuations that smooth the density. The second term describes the effects of flow that translates probability mass. The ergodic density $p := p(\tilde{x} | m) = \mathcal{E}(\Lambda)$ is the principal eigensolution of the Fokker-Planck operator (with an eigenvalue of zero: $\Lambda \mathcal{E} = 0$). Crucially, this density depends only on flow and the amplitude of the controlled fluctuations.

The ergodic density at any point in state-space is also the *sojourn* time that an individual spends there. Similarly, its conditional entropy or ensemble average of *surprise* (also known as self-information or surprisal) is the long-term average of surprise an individual experiences. The entropy and surprise associated with the hidden states are (in the long term: $T \rightarrow \infty$):

$$\begin{aligned}\mathcal{H}(X | m) &= - \int_X p(\tilde{x} | m) \ln p(\tilde{x} | m) dx = \frac{1}{T} \int_0^T dt \mathcal{L}(\tilde{x}(t)) \\ \mathcal{L}(\tilde{x}(t)) &= - \ln p(\tilde{x}(t) | m).\end{aligned}\tag{5}$$

The conditional entropy is an *equivocation* because it is conditioned on the agent. It is important not to confuse the conditional entropy $\mathcal{H}(X | m)$ with the entropy $\mathcal{H}(X)$: A system with low entropy may have a very high conditional entropy unless it occupies states that are characteristic of the agent (because $p(\tilde{x}(t) | m)$ will be persistently small). We will use these characterisations of the ergodic density extensively below and assume that they are all conditional. Readers with a physics background will note that surprise can be regarded as a Lagrangian, with a path-integral $\int dt \mathcal{L}(\tilde{x}(t)) = T \mathcal{H}(X | m)$ that is proportional to entropy. We will call on this equivalence later. In this paper, Lagrangians are negative log-probabilities or surprise.

The terms entropy and surprise are used here in an information theoretic (Shannon) sense. From a thermodynamic perspective, the ergodic density corresponds to a *steady state*, in which (biological) agents are generally far from thermodynamic equilibrium; even though the ensemble density on their macroscopic states (e.g., intracellular concentrations) is stationary. In computational biology, the notion of *non-equilibrium steady state* is central to the study of the homeostatic cellular biochemistry of microscopic states. In this context, the chemical master equation plays the same role as the Fokker-Planck equation above: see [35, 36] for useful introductions and discussion. However, the densities we are concerned with are densities on macroscopic states $\mathcal{O} \subset X$ that ensure the microscopic states they enslave are far from thermodynamic equilibrium. It is these macroscopic states that are characteristic of biological agents. See [37, 38] for useful treatments in the setting of Darwinian dynamics. Having introduced the notion of entropy under ergodic assumptions, we next consider the implications of ergodicity for the flow or motion of agents through their state-space.

2.3. Global Random Attractors. A useful perspective on ergodic agents is provided by the theory of random dynamical systems. A random dynamical system is a measure-theoretic formulation of the solutions to stochastic differential equations like (3). It consists of a base flow (caused by random fluctuations) and a cocycle dynamical system (caused by flow). Ergodicity means the external states constitute a random invariant set $\mathcal{A}(\tilde{\omega}) \subset X$ known as a *pullback* or *global random attractor* [19]. A random attractor can be regarded as the set to which a system evolves after a long period of time (or more precisely the pullback limit, after evolving the system from the distant past to the present: the pullback limit is required because random fluctuations make the system nonautonomous). In the limit of no random fluctuations, random attractors coincide with the definition of a deterministic attractor; as the minimal compact invariant set that attracts all deterministic bounded sets. Crucially, random attractors are compact subsets of state-space that are bounded by deterministic sets. Technically speaking, if the base flow is ergodic and $p(\mathcal{A}(\tilde{\omega}) \subset \mathcal{O}) > 0$ then $\mathcal{A}(\tilde{\omega}) = \Omega_{\mathcal{O}}(\tilde{\omega})$, almost surely [39]. Put simply, this means that if the random attractor falls within a bounded deterministic set $\mathcal{O} \subset X$, then it constitutes an omega limit set $\Omega_{\mathcal{O}}(\tilde{\omega})$. These are the states visited after a sufficiently long period, starting anywhere in $\mathcal{O} \subset X$. In short, if agents are random dynamical systems that spend their time within $\mathcal{O} \subset X$, then they have (are) random attractors.

This existence of random attractors is remarkable because, in the absence of self-organising flow, the fluctuation theorem says they should not exist [40]. The fluctuation theorem generalises the second law of thermodynamics and states that the probability of a system's entropy decreasing vanishes exponentially with time. Put simply, random fluctuations disperse states, so that they leave any bounded set with probability one. See [41] and Appendix A, which show that in the absence of flow

$$\dot{\mathcal{H}}(X | m) = \int_X \frac{\nabla p \cdot \Gamma \cdot \nabla p}{p(\tilde{x} | m)} d\tilde{x} \geq 0. \quad (6)$$

This says that random fluctuations increase entropy production in proportion to their amplitude and the roughness $\nabla p \cdot \nabla p$ of the ensemble density. In the absence of flow, the entropy increases until the density has dispersed and its gradients have been smoothed away. One can think of entropy as the volume or Lebesgue measure $\lambda(\mathcal{A}(\tilde{\omega}))$ of the attracting set: attractors with a small volume concentrate probability mass and reduce average surprise. One can see this heuristically by pretending that all the states within the attractor are visited with equal probability, so that $p(\tilde{x} | m) = 1/\lambda : \tilde{x} \in \mathcal{A}(\tilde{\omega})$. Under this assumption, one can see from (5) that $\mathcal{H}(X | m) = \ln \lambda$ and that entropy increases with volume (and does so more acutely for small volumes). A low entropy means that a small number of states have a high probability of being occupied while the remainder have a low probability. This means that agents with well-defined characteristics have attractors with small measure and an ergodic density with low entropy. The implication here is that agents must counter the dispersive effects of random fluctuations to maintain a high ergodic density over the states

$\mathcal{O} \subset X$ they can occupy. It is important not to confuse the measure of an attracting set $\lambda(\mathcal{A}(\tilde{\omega}))$ with its topological complexity (although, strictly speaking, random attractors are a metric concept not topological). An attractor can have a small measure and yet have a complicated and space-filling shape. Indeed, one might suppose that complex agents (like us) have very complicated random attractors that support diverse and itinerant trajectories; like driving a car within a small margin of error.

2.4. Autopoiesis and Attracting Sets. The formulation of agents as ergodic random dynamical systems has a simple implication: it requires their flow to induce attractors and counter the dispersion of states by random fluctuations. In the absence of this flow, agents would encounter phase-transitions where macroscopic states collapse, exposing their microscopic states to thermodynamic equilibrium. But how do these flows arise? The basic premise, upon which the rest of this paper builds, is that these attractors are autopoietic [42] or self-creating (from the Greek: *auto* ($\alpha\upsilon\tau\acute{o}$) for self- and *poiesis* ($\pi\acute{o}\iota\eta\sigma\iota\varsigma$) for creation). More formally, they arise from the minimisation of entropy with respect to action,

$$a^* = \arg \min_a \mathcal{H}(X | m). \quad (7)$$

Action is the key to creating low entropy densities (resp., low measure attractors), because action determines flow and flow determines the ergodic density (resp., random attractor). This density is the eigensolution $\mathcal{E}(\Lambda(\mathbf{f}, \Gamma))$ of the Fokker-Planck operator, which depends on the policy through the deterministic part of action and the amplitude of random fluctuations through the fluctuating part. This means action plays a dual role in controlling flow to attractive states and suppressing random fluctuations. Equation (6) shows that increasing the amplitude of controlled fluctuations increases the rate of entropy production, because $\partial_f \dot{\mathcal{H}}(X | m) > 0$. This means the fluctuating part of action ω_u can minimise entropy production by suppressing the difference $\tilde{\omega}_x = \tilde{x} - \mathbf{f}(\tilde{x}, u) = \tilde{\omega}_a + \partial_u f \cdot \omega_u + \dots$ between the flow experienced and that expected under the policy. This entails countering unexpected or random deviations from the policy to ensure an autopoietic flow (cf. a ship that maintains its bearing in the face of fluctuating currents and tides). In the absence of fluctuations, flow becomes deterministic and the random attractor becomes a deterministic attractor in the conventional sense (however, it is unlikely that action will have sufficient degrees of freedom to suppress controlled fluctuations completely). Note that for action to suppress random fluctuations about the expected flow (the policy) the agent must have a policy. We will address the emergence and optimisation of policies in the next section. At present, all we are saying is that action must minimise entropy and, implicitly, the measure of an agent's random attractor.

2.5. Summary. In summary, the ergodic or ensemble perspective reduces questions about adaptive behaviour to understanding how motion through state-space minimises surprise and its long-term average (conditional entropy). Action ensures motion conforms to an autopoietic flow or policy,

given the agent and its current state. This policy induces a random invariant set $\mathcal{A}(\tilde{\omega}) \subset \mathcal{O}$ for each class of agent or species, which can be regarded as a probabilistic definition of the agent. This perspective highlights the central role played by the policy: it provides a reference that allows action to counter random fluctuations and violate the fluctuation theorem. In conclusion, the ergodic densities (resp. global random attractors) implied by the existence of biological agents are the stationary solutions to an autopoietic minimisation of their conditional entropy (resp. measure). In the next section, we consider what this implies for the functional anatomy of agents.

3. The Free Energy Formulation

In this section, we introduce the free energy principle as a means of minimising the conditional entropy of an agent's states through action. As noted above, these states and their entropy are hidden from the agent and can only be accessed through sensory states. This means that action cannot minimise the entropy of hidden states directly. However, it can do so indirectly by minimising the entropy of sensory states,

$$a^* = \arg \min_a \mathcal{H}(X | m) = \arg \min_a \mathcal{H}(S | m). \quad (8)$$

This equivalence follows from two assumptions: there is a diffeomorphic mapping between hidden and sensory states and that Jacobian of this mapping (i.e., the sensitivity of sensory signals to their causes) is constant over the range of hidden states encountered (see Appendix B). Crucially, because sensory entropy is the long-term average of sensory surprise, the extremal condition above requires action to minimise the path integral of sensory surprise. This means (by the fundamental lemma of variational calculus) for $t \in [0, T]$

$$\begin{aligned} \delta_a \mathcal{H}(S | m) = 0 &\iff \partial_{a(t)} \mathcal{L}(\tilde{s}(t)) = 0 \iff a(t)^* \\ &= \arg \min_{a(t)} \mathcal{L}(\tilde{s}(t)) \\ \mathcal{H}(S | m) &= \frac{1}{T} \int_0^T dt \mathcal{L}(\tilde{s}(t)) \\ \mathcal{L}(\tilde{s}(t)) &= -\ln p(\tilde{s}(t) | m). \end{aligned} \quad (9)$$

Equation (9) says that it is sufficient for action to minimise sensory surprise to minimise the entropy of sensations (or at least find a local minimum). This is sensible because action should counter surprising deviations from the expected flow of states. However, there is a problem; agents cannot evaluate sensory surprise $\mathcal{L}(\tilde{s}(t))$ explicitly, because this would involve integrating $p(\tilde{s}, \tilde{x}, \theta | m)$ over hidden states and parameters or causes: $\vartheta = (\tilde{x}, \theta)$. This is where the free energy comes in.

Free energy is a functional of sensory and internal states that upper bounds sensory surprise and can be minimised through action (cf. (2)). Effectively, free energy allows agents to finesse a generally intractable integration problem (evaluating surprise) by reformulating it as an optimisation problem. This well-known device was introduced by Feynman [43] and has been exploited extensively in machine

learning and statistics [44–46]. The requisite free energy bound is created by adding a nonnegative Kullback-Leibler divergence or cross-entropy term [47] to surprise:

$$\begin{aligned} \mathcal{F}(t) &= \mathcal{L}(\tilde{s}(t)) + D_{\text{KL}}(q(\vartheta) || p(\vartheta | \tilde{s}, m)) \\ &= \langle \ln q(\vartheta) \rangle_q - \langle \ln p(\tilde{s}, \vartheta | m) \rangle_q. \end{aligned} \quad (10)$$

This divergence is induced by a recognition density $q(\vartheta) := q(\vartheta | \tilde{\mu})$ on the hidden causes of sensory states. This density is associated with the agent's internal states $\tilde{\mu}(t)$ that play the role of sufficient statistics; for example, the mean or expectation of hidden causes. Free energy $\mathcal{F}(\tilde{s}, \tilde{\mu}) \in \mathbb{R}$ can be evaluated because it is a functional of internal states and a generative model $p(\tilde{s}, \vartheta | m)$ entailed by the agent. This can be seen from second equality, which expresses free energy in terms of the negentropy of $q(\vartheta)$ and the expected value of $\ln p(\tilde{s}, \vartheta | m)$.

To ensure action minimises surprise, the free energy must be minimised with respect the internal variables that encode the recognition density (to ensure the free energy is a tight bound on surprise). This is effectively perception because the cross-entropy term in (10) is non-negative, with equality when $q(\vartheta | \tilde{\mu}) = p(\vartheta | \tilde{s}, m)$ is the true conditional density. In short, optimising the recognition density makes it an approximate conditional density on the causes of sensory states. This is the basis of perceptual inference and learning as articulated by the Bayesian brain hypothesis [10, 13, 48–52]. We can now formulate action (9) in terms of a dual minimisation of free energy (see (2) and Figure 1).

$$\begin{aligned} a^* &= \arg \min_a \mathcal{F}(\tilde{s}, \tilde{\mu}), \\ \tilde{\mu}^* &= \arg \min_{\tilde{\mu}} \mathcal{F}(\tilde{s}, \tilde{\mu}). \end{aligned} \quad (11)$$

Action minimises free energy through changing the generalised motion of hidden states. In essence, it ensures that the trajectory of sensory states conform to the agents conditional beliefs encoded by internal states. Note that action is fundamentally different from a policy in optimal control and reinforcement-learning. Action is not a deterministic function of hidden states and is sensitive to random fluctuation causing sensory states. This means, unlike an optimal policy, it can suppress surprises by countering unexpected fluctuations in sensory states: although optimal control schemes can recover from perturbations, they cannot cancel them actively. However, as we will see below, optimal policies play a key role providing in prior constraints on the flow of hidden states that action tries to disclose.

3.1. Active Inference and Generalised Filtering. In what follows, we will assume that the minimisation of free energy with respect to action and internal states (11) conforms to a generalised gradient descent,

$$\begin{aligned} \dot{a} &= -\partial_a \mathcal{F}(\tilde{s}, \tilde{\mu}), \\ \dot{\tilde{\mu}} &= \mathcal{D}\tilde{\mu} - \partial_{\tilde{\mu}} \mathcal{F}(\tilde{s}, \tilde{\mu}). \end{aligned} \quad (12)$$

These coupled differential equations describe action and perception respectively. The first just says that action suppresses free energy. The second is known as generalised filtering [53] and has the same form as Bayesian (e.g., Kalman-Bucy) filtering, used in time series analysis. The first term is a prediction based upon the differential operator \mathcal{D} that returns the generalised motion of internal states encoding conditional predictions. The second term is usually expressed as a mixture of *prediction errors* that ensures the internal states (sufficient statistics) are updated in a Bayes-optimal fashion (see below). The differential equations above are coupled because sensory states depend upon action, which depends upon perception through the conditional predictions. This circular dependency leads to a sampling of sensory input that is both predicted and predictable, thereby minimising free energy and surprise. This is known as active inference.

In generalised filtering, one treats hidden parameters as hidden states that change very slowly: the ensuing generalised descent can then be written as a second-order differential equation: $\ddot{\mu}_\theta = -\partial_\theta \mathcal{F} - \kappa \mu'_\theta$, where κ is the (high) prior precision on changes in hidden parameters. See [53] for details. In neurobiological formulations of free energy minimisation, internal states generally correspond to conditional expectations about hidden states and parameters, which are associated with neuronal activity and connections strengths, respectively. In this setting, optimising the conditional expectations about hidden states (neuronal activity) corresponds to *perceptual inference* while optimising conditional expectations about hidden parameters (neuronal plasticity) corresponds to *perceptual learning*.

Equation (12) describes the dynamics of action and internal states, whose particular form depends upon the generative model of the world. We will assume this model has the following form:

$$\begin{aligned}\dot{\tilde{x}} &= f(\tilde{x}, \theta) + \tilde{\omega}_x, \\ \tilde{s}(t) &= g(\tilde{x}, \theta) + \tilde{\omega}_s.\end{aligned}\quad (13)$$

As in the previous section, (f, g) are nonlinear functions of hidden states that generate sensory states; however, these are distinct from the real equations of motion and sensory mappings (\mathbf{f}, \mathbf{g}) that depend on action. The generative model does not include action, because action is not a hidden state. Random fluctuations (ω_s, ω_x) play the role of sensory noise and induce uncertainty about the motion of hidden states. Hidden states are abstract quantities (like the motion of an object in the field of view) that the agent uses to explain or predict sensations. Gaussian assumptions about the random fluctuations in (13) furnish a probabilistic generative model of sensory states $p(\tilde{s}, \vartheta \mid m)$ that is necessary to evaluate free energy. See [53] for a full description of generalised filtering in the context of hierarchical dynamic models. For simplicity, we have assumed that state-space associated with the generative model is the same as the hidden state-space in the world. However, this is not necessary, because exchanges with the environment are mediated through sensory states and action.

Given the form of the generative model (13) and an assumed (Gaussian) form for the recognition density, we can now write down the differential equations (12) describing the

dynamics of internal states in terms of (precision-weighted) prediction errors $(\tilde{\varepsilon}_s, \tilde{\varepsilon}_x)$ on sensory states and the predicted motion of hidden states, where (ignoring some second-order terms and using $\tilde{g} := g(\tilde{x}, \theta)$)

$$\begin{aligned}\dot{\tilde{\mu}} &= \mathcal{D}\tilde{\mu} + \partial_{\tilde{\mu}}\tilde{g} \cdot \tilde{\varepsilon}_s + \partial_{\tilde{\mu}}\tilde{f} \cdot \tilde{\varepsilon}_x - \mathcal{D}^T\tilde{\varepsilon}_x, \\ \tilde{\varepsilon}_s &= \tilde{\Pi}_s(\tilde{s} - \tilde{g}), \\ \tilde{\varepsilon}_x &= \tilde{\Pi}_x(\mathcal{D}\tilde{\mu} - \tilde{f}).\end{aligned}\quad (14)$$

The (inverse) amplitude of generalised random fluctuations are encoded by their precision $(\tilde{\Pi}_s, \tilde{\Pi}_x)$, which we assume to be fixed in this paper. This particular free energy minimisation scheme is known as *generalised predictive coding* and has become a useful metaphor for neuronal message passing in the brain: see also [12]. The simplicity of this scheme stems from the assumed Gaussian form of the recognition density. This means the internal states or sufficient statistics can be reduced to conditional expectations (see Appendix C). In neural network terms, (14) says that error-units receive predictions while prediction-units are driven by prediction errors. In neurobiological implementations of this scheme, the sources of prediction errors are usually thought to be superficial pyramidal cells while predictions are conveyed from deep pyramidal cells to superficial pyramidal cells encoding prediction error [54].

Because action can only affect the free energy by changing sensory states, it can only affect sensory prediction errors. From (13), we have

$$\dot{a} = -\partial_a \tilde{s} \cdot \tilde{\varepsilon}_s.\quad (15)$$

In biologically plausible instances of active inference, the partial derivatives in (15) would have to be computed on the basis of a mapping from action to sensory consequences, which is usually quite simple; for example, activating an intrafusal muscle fibre elicits stretch receptor activity in the corresponding spindle: see [6] for discussion.

3.2. Summary. In summary, we can account for the unnatural persistence of self-organising biological systems in terms of action that counters the dispersion of their states by random fluctuations. This action minimises the entropy of their ergodic density by minimising a free energy bound on sensory surprise or self-information as each point in time. To ensure the free energy is a good proxy for surprise, internal states must also minimise free energy and implicitly represent hidden states. This minimisation rests upon a generative model, which furnishes conditional predictions that action can fulfil. These predictions rest on equations of motion that constitute (empirical) priors [55] on the flow of hidden states in the world. In short, agents are equipped with a model of dynamics in their local environment and navigate that environment to minimise their surprise.

We can now associate the expected flow of the previous section with the empirical priors learned under the generative model: $\mathbf{f}(x, u) = f(x, \mu_\theta)$. This rests upon the assumption that action eliminates (on average) the difference between the actual and predicted flow. This means the predicted

flow corresponds to the policy. The policy $\mathbf{f}(x, u) = f(x, \mu_\theta)$ is an *empirical* prior because it depends on conditional beliefs about hidden parameters encoded by μ_θ . This is an important point because it means that the environment causes prior beliefs about motion (through parameter learning), while these beliefs cause the sampled environment. This circular causality is the basis of autopoietic flow and highlights the fact self-organisation rests on a reciprocal exchange between implicit beliefs about how an agent or system will behave and behavioural constraints that are learned by behaving. Minimising free energy ensures that the beliefs and constraints are consistent and enables the agent to create its own environment. In this view, perceptual inference becomes truly embodied or situated and is an integral part of sustainable interactions with the environment. The previous section suggested that action was the key to understanding self-organised behaviour. This section suggests that action depends on a policy or empirical priors over flow. In what follows, we consider the nature of this flow and its specifications.

4. Policies and Value

The previous section established differential equations that correspond to action and perception under a model of how hidden states evolve. These equations are based on the assumption that agents suppress (a bound on) surprise and, implicitly, the entropy of their ergodic density. We now consider optimising the model per se, in terms of formal priors on flow. These correspond to the form of the equation of motions in (13). In particular, we will consider constraints encoded by a (cost) function $c(\tilde{x}) \subset m$ over hidden states. The existence of autopoietic flow is not mysterious, in the sense that agents who do not have a random attractor cannot exist. In other words, every agent (phenotype) can be regarded as a solution to the Fokker-Planck equation, whose policy is compatible with the biophysics of its environmental niche. One might conjecture that each solution (random attractor) corresponds to a different species, and that there may be a limited number of solutions as evidenced by convergent evolution [17]. This section considers the policies that underwrite these solutions and introduces the notion of *value* in terms of the Helmholtz decomposition. In brief, we will see that flow determines value, where value is negative surprise.

We start with the well-known decomposition of flow into curl- and divergence-free components (strictly speaking, the first term is only curl-free when $\Gamma(\tilde{x}) = \gamma(\tilde{x}) \cdot I$; that is, the diffusion tensor is isotropic. However, this does not affect the following arguments, which rest on the divergence-free component),

$$f = \Gamma \cdot \nabla V + \nabla \times W. \quad (16)$$

This is the Helmholtz decomposition (also known as the fundamental theorem of vector calculus) and expresses any policy in terms of scalar $V(\tilde{x})$ and vector $W(\tilde{x})$ potentials that prescribe irrotational (curl-free) $\Gamma \cdot \nabla V$ and solenoidal (divergence-free) $\nabla \times W$ flow. An important decomposition described in [37, 56], formulates the divergence-free part in

terms of an antisymmetric matrix, $Q(\tilde{x}) = -Q(\tilde{x})^T$ and the scalar potential, which we will call *value*, such that

$$\begin{aligned} f &= (\Gamma + Q)\nabla V \implies \\ \nabla \times W &= Q\nabla V. \end{aligned} \quad (17)$$

Using this (*standard form*) decomposition [57], it is fairly easy to show that $p(\tilde{x} | m) = \exp(V(\tilde{x}))$ is the equilibrium solution to the Fokker-Planck equation (4):

$$\begin{aligned} p &= \exp(V) \implies \nabla p = p\nabla V \implies \\ \Lambda p &= \nabla \cdot \Gamma \nabla p - \nabla \cdot (fp) \\ &= -p(\nabla \cdot (Q\nabla V) + (Q\nabla V) \cdot \nabla V) = 0. \end{aligned} \quad (18)$$

Equation (18) uses the fact that the divergence-free component is orthogonal to ∇V (see Appendix D). This straightforward but fundamental result means that the flow of any ergodic random dynamical system can be expressed in terms of orthogonal curl- and divergence-free components, where the (dissipative) curl-free part increases value while the (conservative) divergence-free part follows isoprobability contours and does not change value. Crucially, under this decomposition value is simply negative surprise: $\ln p(\tilde{x} | m) = V(\tilde{x}) = -\mathcal{L}(\tilde{x} | m)$. It is easy to show that surprise (or value) is a Lyapunov function for the policy

$$\begin{aligned} \dot{V}(x(t)) &= \nabla V \cdot f = \nabla V \cdot \Gamma \cdot \nabla V + \nabla V \cdot \nabla \times W \\ &= \nabla V \cdot \Gamma \cdot \nabla V \geq 0. \end{aligned} \quad (19)$$

Lyapunov functions always decrease (or increase) with time and are used to establish the stability of fixed points in deterministic dynamical systems. This means every policy (expected flow) reduces surprise as a function of time. In other words, it must direct flow towards states that are more probable (and have a greater sojourn time). This is just a formal statement of the fact that ergodic systems must, on average, continuously suppress surprise, to offset the dispersive effect of random fluctuations. Ao reviews the importance and generality of the decomposition in (17) and how it provides a unifying perspective on evolutionary and statistical dynamics [38]: this decomposition shows that fluctuations in Darwinian dynamics imply the existence of canonical distributions of the Boltzmann-Gibbs type. Furthermore, it demonstrates the second law of thermodynamics, without detailed balance. In particular, the dynamical (divergence-free) component responsible for breaking detailed balance does not contribute to changes in entropy. In short, (17) represents “a simple starting point for statistical mechanics and thermodynamics and is consistent with conservative dynamics that dominates the physical sciences” [58]. The generality of this formulation can be appreciated by considering two extreme cases of flow that emphasise the curl and divergence-free components, respectively.

4.1. Conservative (Divergence-Free) Flow. When the random fluctuations are negligible (i.e., $\Gamma \rightarrow 0$), irrotational (curl-free) flow $\Gamma \cdot \nabla V = 0$ disappears and we are left with

divergence-free flow that describes conservative dynamics (e.g., classical mechanics). These flows would be appropriate for massive bodies with virtually no random fluctuations in their motion. A simple example would be the Newtonian mechanics that result from a Lagrangian (surprise) and anti-symmetric matrix,

$$\mathcal{L}(\tilde{x}) = \varphi(x) + \frac{1}{2}x'^2$$

$$Q = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \Rightarrow f = -Q\nabla\mathcal{L} = \begin{bmatrix} \dot{x} \\ \dot{x}' \end{bmatrix} = \begin{bmatrix} x' \\ -\nabla\varphi \end{bmatrix}. \quad (20)$$

This describes the motion of a unit mass in a potential field $\varphi(x)$, where the Lagrangian comprises potential and kinetic terms. Things get more interesting when we consider random fluctuations in the velocity,

$$\Gamma = \begin{bmatrix} 0 & 0 \\ 0 & \gamma \end{bmatrix} \Rightarrow$$

$$f = -(\Gamma + Q)\nabla\mathcal{L} = \begin{bmatrix} x' \\ -\nabla\varphi - \gamma x' \end{bmatrix}. \quad (21)$$

This introduces a motion-dependent reduction in the motion of velocity (acceleration) that corresponds to friction. Note that friction is an emergent property of random fluctuations in velocity (and nothing more). A more thorough treatment of the relationship between the diffusion due to random fluctuations and friction can be found in [57], using the generalised Einstein relation. Consider now systems in which random fluctuations dominate and the conservative (divergence-free) flow can be ignored.

4.2. Dissipative (Curl-Free) Flow and Detailed Balance. Here, irrotational (curl-free) flow dominates and the dynamics have *detailed balance*, which means that flow can be expressed as an ascent on a scalar (value) potential: $f = \Gamma \cdot \nabla V = -\Gamma \cdot \nabla\mathcal{L}$. Crucially, because there is effectively no conservative flow, the ergodic density concentrates around the maximum of value (or minimum of surprise), which (in the deterministic limit) induces a fixed point attractor. Curl-free policies are introduced here, because of their central role in optimal control and decision (game) theory: in the next section, we will consider curl-free policies that are specified in terms of value-functions, $V(\tilde{x})$. These range from reinforcement-learning heuristics in psychology to more formal optimal control theory treatments. However, one should note that these approaches are incomplete in the sense they do not specify generic policies: a complete specification of flow would require the vector potential $W(\tilde{x})$ or, equivalently, the anti-symmetric matrix, $Q(\tilde{x})$. This means that it is not sufficient to know (or learn) the value of a state to specify a policy explicitly, unless the environment permits curl-free policies with detailed balance (i.e., with no classical or conservative dynamics).

Ergodic densities under detailed balance are closely connected to *quantal response equilibria* (QRE) in economics and game theory. QRE are game-theoretical formulations that

provide an alternative to Nash equilibria [18]. QRE do not require perfect rationality; players are assumed to make normally distributed errors in their predicted payoff. In the limit of no errors, QRE predict unique Nash equilibria. From the point of view of game theory, the interesting questions pertain to different equilibria prescribed by the policy or state-transitions. These equilibria are analogous to the solutions of the Fokker-Planck equation above, where $V(\tilde{x})$ is called *attraction* and $\Gamma \in \mathbb{R}^+$ is temperature or inverse sensitivity [9, 59]. In this context, the ergodic density $p(\tilde{x} | m) = \exp(-\mathcal{L})$ prescribes optimal states or choices probabilistically, in terms of value, where $V = -\mathcal{L}$. This prescription is closely related to softmax or logit discrete choice models [60], which are the most common specification of QRE. In economics, optimal state-transitions lead to equilibria that maximise value or expected utility. These are low-entropy densities with probability mass on states with high utility. We pursue this theme in below, in the context of optimal control theory and reinforcement-learning.

4.3. Summary. In this section, we have seen that a policy or empirical priors on flow (specified by conditional beliefs about the parameters of equations of motion) can be decomposed into curl and divergence-free components, specified in terms of a value-function and antisymmetric matrix that determines conservative flows of the sort seen in classical mechanics. Crucially, this value-function is just negative surprise and defines the ergodic (invariant) probability density over hidden states. However, we have no idea about where the policy comes from. All we know is that it furnishes a solution to the Fokker-Planck equation; an idiocentric description of an agent's exchange with the environment. The remainder of this paper will be concerned with how policies are specified and how they are instantiated in terms of value-functions.

Evolutionary theory [61, 62] suggests that species (random attractors) do not arise *de novo* but evolve through natural selection (e.g., by punctuated equilibria or phyletic gradualism; [63, 64]). We take this to imply that policies are heritable and can be encoded (epigenetically) in terms of value or cost-functions. We will assume the agents are equipped with a cost-function that labels states as attractive or not

$$c(x | m) \leq 0 : \quad x \in \mathcal{A} = \bigcap_{\tilde{\omega} \in \Omega} \mathcal{A}(\tilde{\omega}), \quad (22)$$

$$c(x | m) > 0 : \quad x \notin \mathcal{A}.$$

Technically, cost indicates whether each state is in a *kernel* or the set of *fixed points* of a random attractor [65]. In the deterministic limit $\Gamma \rightarrow 0$ this kernel reduces to an attractor in the usual sense. From now on, we will use $\mathcal{A} \subset \mathcal{O}$ to mean the kernel of a random attractor or an attractor in the deterministic sense. The introduction of cost allows us to connect attractors in dynamical systems with attractive states in reinforcement-learning and optimal control. Informally, cost labels states as either attractive (e.g. sated) or costly (e.g., thirsty). The cost-function could also be regarded as a characteristic function that indicates whether the current state is characteristic of the class the agent belongs to. This labelling is sufficient to prescribe policies that assure

equilibrium solutions, as evidenced by the existence of evolved agents. This question is how? We will begin by considering control theory.

5. Optimal (Fixed Point) Control and Reinforcement-Learning

In this section, we look at policies and value from the point of view of optimal control theory and reinforcement-learning. In the previous section, value was considered to arise from a decomposition of flow into curl, and divergence-free parts. In that setting, value simply reports the surprise that a state is occupied. In other words, value is an attribute of the policy. Optimal control theory turns this around and assumes that the policy is an attribute of value. This enables policies to be specified by value, via cost-functions. In this section, we will consider optimal control theory as optimising policies (flows), whose attracting fixed-points are specified by cost-functions. Crucially, because optimal control policies do not specify divergence-free flow, they can only specify policies with attracting fixed points (the maxima of the value function). In the next section, we turn to generalised policies that exploit divergence-free flow to support itinerant policies. We will persist with continuous time formulations in this section and provide discrete time versions of the main results in the appendices.

5.1. Optimal Control Theory. In optimal control theory and its ethological variants (i.e., reinforcement-learning), adaptive behaviour is formulated in terms of how agents navigate state-space to access sparse rewards and avoid costly regimes. The aim is to find a (proximal) policy that attains long-term (distal) rewards. In terms of the previous section, a policy $f = \Gamma \cdot \nabla V + \nabla \times W$ was specified via the scalar potential or value $V(\tilde{x})$ also known as (negative) *cost-to-go*. In optimal control theory, value is defined as the expected path-integral of cost. More formally, the cost-to-go of $\tilde{x}_0 \in X$ is the cost expected over future times $t \in [t_0, \infty]$, starting with a point density $p(\tilde{x}, t_0 | m) = \delta(\tilde{x}_0)$, which evolves according to (3) (see Appendix E),

$$V(\tilde{x}_0) = - \int_{t_0}^{\infty} \int c(\tilde{x}) p(\tilde{x}, t | m) dx dt \implies \quad (23)$$

$$c(\tilde{x}) = f \cdot \nabla V(\tilde{x}) + \nabla \cdot \Gamma \cdot \nabla V(\tilde{x}).$$

Or in the deterministic limit $\Gamma \rightarrow 0$

$$V(\tilde{x}_0) = - \int_{t_0}^{\infty} c(\tilde{x}(t)) dt \implies \quad (24)$$

$$c(\tilde{x}) = f \cdot \nabla V(\tilde{x}) = \dot{V}(\tilde{x}(t)).$$

This definition of value as an expected path-integral of cost (first line) allows cost to be expressed as a function of value (second line). It says that cost is the expected increase in value. This may sound counterintuitive but makes sense if one considers a reward now means less in the future (i.e., a decrease in the value of the next state). Crucially, (24) shows that the maxima of the ergodic density can only exist where

cost is zero or less (cf. (22)): at a maximum of $p(\tilde{x} | m) = \exp(V)$ we have the following:

$$\left. \begin{array}{l} \nabla V(\tilde{x}) = 0 \\ \nabla \cdot \Gamma \cdot \nabla V(\tilde{x}) \leq 0 \end{array} \right\} \implies c(\tilde{x}) \leq 0, \quad (25)$$

with $c(\tilde{x}) = 0$ in the deterministic limit. Put simply, costly regions induce value gradients that guide flow towards points where there is no cost (i.e., no gradients). In this sense, value is sometimes called a navigation function. This means that, in principle, we have a way to prescribe ergodic densities with maxima (attracting fixed-points) that are specified with a cost-function. Equation (23) shows that the cost-function can be derived easily, given the policy and implicit value-function. However, to specify a policy with cost, we have to derive the flow from the cost-function. This entails associating a unique flow with the value-function and solving (24) for value: this association is necessary because optimal control does not specify the divergence-free part of the policy. Solving (24) is the difficult problem optimal control and value-learning deal with.

Let optimal control be denoted by $\pi(\tilde{x})$, where optimal control maximises the ascent of the value-function

$$\pi(\tilde{x}) = \arg \max_u f(\tilde{x}, u) \cdot \nabla V(\tilde{x}). \quad (26)$$

This extremal condition associates a unique (optimal) flow with every value-function such that value can be expressed in terms of cost using (24)

$$\max_u f(\tilde{x}, u) \cdot \nabla V(\tilde{x}) = f(\tilde{x}, \pi) \cdot \nabla V(\tilde{x}) = c(\tilde{x}). \quad (27)$$

This is the celebrated Hamilton-Jacobi-Bellman (HJB) equation. More general forms are provided in Appendix F for the interested reader. The basic problem, posed by the solution of the HJB equation for value, is that the value-function depends on optimal control, so that future cost can be evaluated. However, optimal control depends on the value-function. This circular dependency can only be resolved by solving the self-consistent equations above, also known as the dynamic programming recurrence. This is the *raison d'être* for value-learning.

In engineering, planning, and control problems, the HJB equation is usually solved by backwards induction (staring at the desired fixed-point and working backwards). However, this is not an ethological option for agents that have to learn the value-function online. An alternative solution exploits the fact that the expected increase in the value of the current state is cost. This leads to a straightforward value-learning scheme

$$V(x(t)) \leftarrow V(x(t)) + \delta : \delta = \dot{V}(x(t)) - c(x(t)). \quad (28)$$

Such that, at convergence, the value-function satisfies the deterministic limit of (24), at least for the states visited,

$$\delta(t) \rightarrow 0 \implies c(x) = \dot{V}(x(t)) = f \cdot \nabla V(x). \quad (29)$$

Heuristically, this scheme erodes the value landscape, creating gradients that ensure flow through costly regions of state-space. The only points that are exempt from this erosion are

maxima with zero flow and cost (or negative cost in the presence of fluctuations). These are the fixed-points of the attracting set prescribed by the cost-function.

In (28), $\delta(t)$ reports the difference between the cost predicted $\hat{V}(x(t))$ and observed $c(x(t))$. This is the negative cost or reward prediction error. There is a vast literature on reinforcement-learning schemes that solve the discrete time version of the HJB equation; either by backwards induction (model-based schemes) or by using reward prediction error (model-free schemes). Model-free schemes use a discrete time version of (28) using the Robbins-Monro algorithm; [66, 67]. Appendix G provides a brief survey of these schemes. Intuitively, they all involve increasing the value of the last state in proportion to the reward prediction error. This is the basis of temporal difference schemes [2] and Rescorla-Wagner [68] models of conditioning in psychology. See [5] for a comprehensive review. If the agent has no model of flow or state-transitions, similar schemes can be invoked to optimise the policy, (e.g., actor-critic schemes). A generalisation of value-learning, called Q-learning, considers a value or *quality* $Q : X \times U \rightarrow \mathbb{R}$ on the joint space of states and control [69]. Q-learning does not need a model in the form of probabilistic transitions to optimise control, because the quality of an action is encoded explicitly. Perhaps the most important thing to come out of these modelling initiatives is that phasic dopamine discharges in the brain are a prime candidate for reporting reward prediction error [3, 70, 71]. In some cases theoretical predictions preempted empirical findings; for example, “in the absence of an expected reward, value system responses should show a decrease in activity at the time when the reward would normally be delivered” [72], which was confirmed subsequently [73].

5.2. Summary. In summary, one can decompose any policy or expected flow into a part that is divergence-free and a part that increases value, where value is negative surprise or the log probability of a state being occupied. This means, given expected flow and the amplitude of random fluctuations about that flow, one can compute the ergodic density and associated value-function. Furthermore, if one defines surprise (negative value) of any state as the expected cost accumulated from that state, then it is straightforward to evaluate the cost-functions implied by any given policy. An example is given in Figure 4 using the Lorentz attractor in previous figures.

Using this definition of cost, reinforcement-learning and optimal control theory try to derive value from cost-functions by assuming controlled flow minimises accumulated cost. This involves solving self-consistent equations that entail these assumptions. The ensuing value-function guides flow to ensure cost is minimized under constraints on motion. In dynamical terms, this approach optimises a policy in terms of its scalar potential, whose maxima coincide with points in state-space where $c(x) \leq 0$. However, despite its prominence in the neurosciences, optimal control theory is not very useful for understanding self-organisation in biological systems. For example, one could not specify the policy followed by the Lorentz attractor in Figure 4, because its dynamics do not have detailed balance. In other words,

although one can derive a cost-function from any flow, one cannot specify any flow with a cost-function: to specify any given policy one would need the vector potentials (or anti-symmetric matrices) above.

Furthermore, optimal control schemes and related heuristics have several shortcomings: (i) they are notoriously slow, requiring hundreds if not thousands of trials to learn even the simplest value-function, (ii) value-learning based on stochastic iteration depends on the same random fluctuations that need to be suppressed to pursue the policy, (iii) optimal control theory says nothing about exploration of state-space, (iv) an exigent limitation of these schemes is that they only account for policies with stationary fixed-points (i.e., agents who would optimally do nothing). This means they cannot account for the itinerant and context-sensitive nature of real behaviour. To resolve these problems we now turn to generalised policies that include divergence-free flow and are constrained directly, as opposed to placing constraints on value functions.

6. Generalised (Itinerant) Policies

In the previous section, cost was used to specify policies or expected flow in terms of value-functions. However, policies with detailed balance (of the form $f = \Gamma \cdot \nabla V$) place severe constraints on the attractors they engender. In the deterministic limit, they can only prescribe attractors with a single fixed-point (per basin of attraction), which is a local maximum of value $V : X \rightarrow \mathbb{R}$. This is not a useful description of real agents that exhibit itinerant dynamics with quasiperiodic and chaotic behaviour. In what follows, we consider policies in generalised coordinates of motion that do not have detailed balance (and exploit divergence-free flow); we will refer to these as generalised policies. These policies provide a rich repertoire of dynamics with attracting sets that can be specified directly by policies. In other words, we dispense with the assumptions of optimal control theory and consider the more general problem of how to specify attracting sets. In what follows, cost-functions are used to specify (unstable) fixed points for fairly generic flows. We will see how this leads to a dynamical formulation of policies and the emergence of optimal itinerancy. This section focuses on the basics and provides some simple examples, noting that there are many potential extensions; some of which have already been addressed in the application of dynamical systems theory to biological systems (see below).

Generalised policies rest on the ensemble dynamics perspective: we start by considering how cost can restrict the probability mass of an ergodic density to a subset of state-space $\mathcal{A} \subset X$. The Fokker-Planck equation (3) provides a fundamental constraint on flow that must be satisfied when $\Delta p = 0$ and

$$\begin{aligned} \nabla \cdot \Gamma \nabla p &= \nabla \cdot (fp) = f \cdot \nabla p + p \nabla \cdot f \implies \\ p(\tilde{x} | m) &= \frac{\nabla \cdot \Gamma \nabla p - f \cdot \nabla p}{\nabla \cdot f}. \end{aligned} \quad (30)$$

This straightforward result shows that as divergence $\nabla \cdot f$ decreases, the sojourn time (i.e., the proportion of time

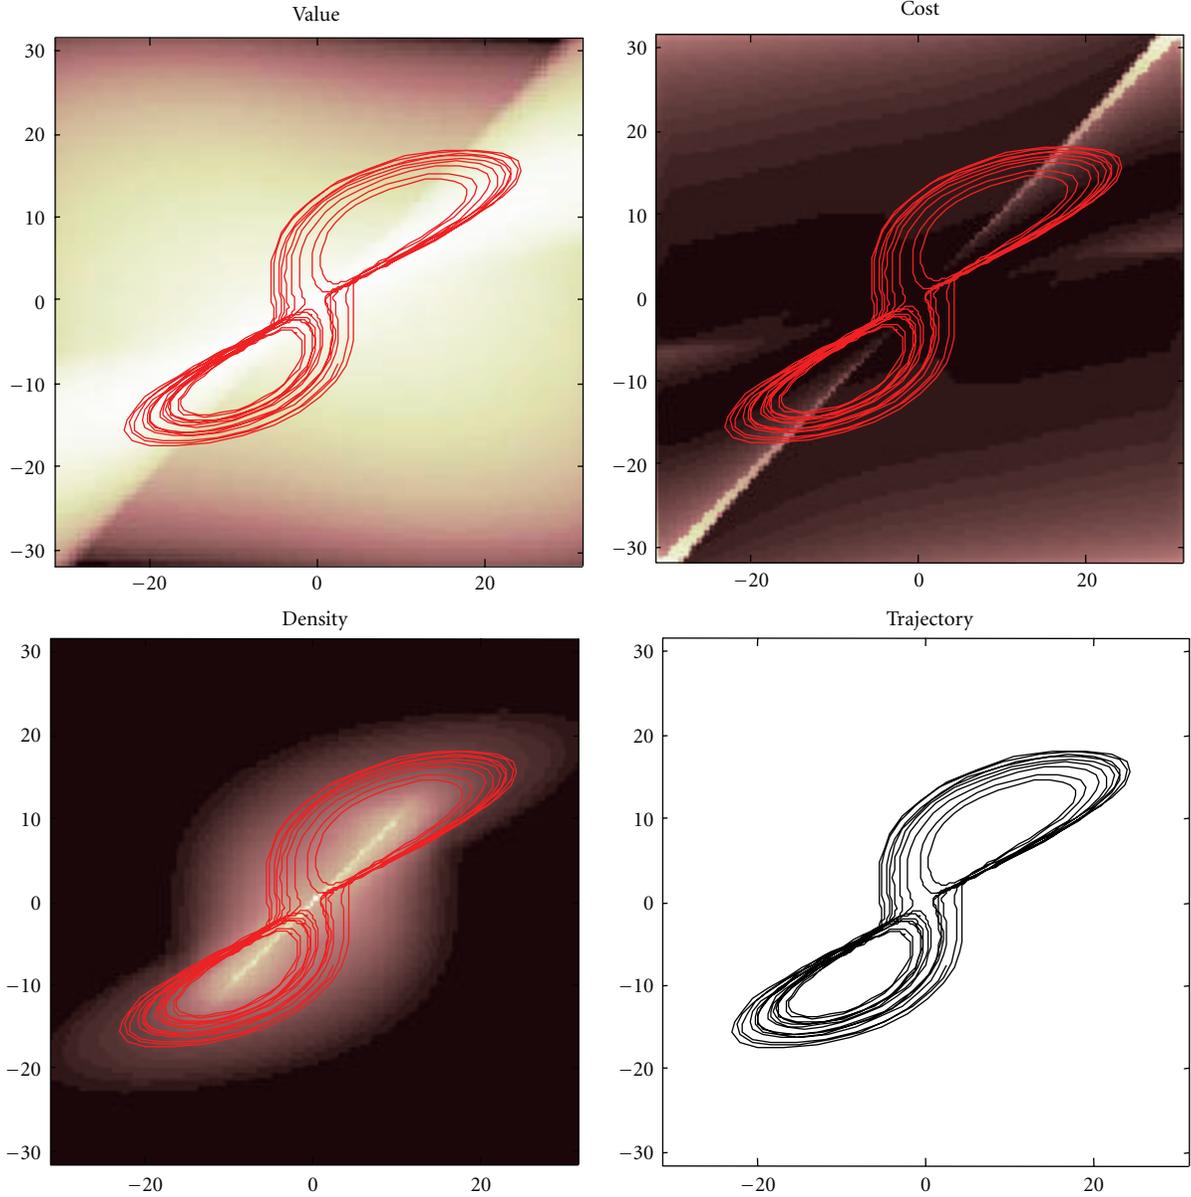


FIGURE 4: Value and cost functions of dynamical systems. This figure shows the value and cost functions of the Lorenz attractor used in the previous figures. These functions always exist for any global random attractor because value (negative surprise) is the log density of the eigensolution of the systems Fokker-Planck operator. This means, given any deterministic motion (flow) and the amplitude of random fluctuations (diffusion), we can compute the Fokker Planck operator $\Lambda(f, \Gamma)$ and its eigensolution $p = \mathcal{E}(\Lambda)$ and thereby define value $V = \ln p$. Having defined value, cost is just the expected rate of change of value, which is given by the deterministic flow and diffusion (see (23)). In this example, we computed the eigensolution or ergodic density using a discretisation of state-space into 96 bins over the ranges: $[-32, 32] \times [-32, 32] \times [4, 64]$ and a diffusion tensor of $\Gamma = (1/64) \cdot I$. The upper panels show the resulting value and (negative) cost functions for a slice through state-space at $x_3 = 24$. Note how cost takes large values when the trajectory (red line) passes through large value gradients. The lower left panel shows the resulting ergodic density as a maximum intensity projection over the third state. A segment of the trajectory producing this density is shown on the lower right.

a state is occupied) rises. This means divergence decreases value and increases surprise. This is intuitive, in that divergence represents the extent to which flow behaves like a source or a sink at a given point. Attractive or valuable points in state-space are sinks with low divergence. At the peaks of the ergodic density its gradient is zero and its curvature is negative:

$$\left. \begin{array}{l} p > 0 \\ \nabla p = 0 \\ \nabla \cdot \Gamma \cdot \nabla p < 0 \end{array} \right\} \Rightarrow \nabla \cdot f < 0. \quad (31)$$

This means divergence must be negative. This generalises an almost trivial result in deterministic systems: divergence is

the sum of the flow’s real Lyapunov exponents $\nabla \cdot f = \text{tr}(\partial_{\tilde{x}} f) = \sum_i \text{Re}(\lambda_i)$, where λ_i are the eigenvalues of the Jacobian $\partial_{\tilde{x}} f$. If divergence is negative, then all the Lyapunov exponents are negative, implying a stable fixed-point attractor.

This provides a straightforward way of ensuring the peaks of the ergodic density lie in, and only in $\mathcal{A} \subset X$. This is assured if $\nabla \cdot f < 0$ when $\tilde{x} \in \mathcal{A}$ and $\nabla \cdot f \geq 0$ otherwise. We can exploit this using the equations of motion in (21)

$$f(x, x') = \left[\begin{array}{c} x' \\ c(x) \cdot x' - \partial_x \varphi \end{array} \right] \implies \nabla \cdot f = c. \quad (32)$$

This flow describes the classical (Newtonian) motion of a unit mass in a potential energy well $\varphi(x)$, where cost plays the role of negative dissipation or friction. Crucially, under this policy or flow, divergence is cost. This means the associated ensemble density can only have maxima in regions of negative (*divergence-based*) cost. This provides a means to specify attractive regions $\mathcal{A} \subset X$ by assigning them negative cost, which brings us back to (22),

$$\begin{aligned} c(x) \leq 0 &: x \in \mathcal{A}, \\ c(x) > 0 &: x \notin \mathcal{A}. \end{aligned} \quad (33)$$

Put simply, this scheme ensures that agents are expelled from high-cost regions of state-space and get “stuck” in attractive regions. Equivalently, cost can be regarded as destroying unattractive fixed points at the minima of the potential landscape $\varphi(x)$. It should be noted that negative divergence does not ensure an attractive fixed point; however, if divergence is sufficiently negative, the point will act as a sink and become, almost surely, part of the attracting set. We will see an example of this below. In summary, in attractive regions with low cost, flow slows down sufficiently to increase sojourn time and implicitly value. This is not dissimilar to Win-Stay, Lose-Shift strategies used to model optimal decisions in game theory [74]. We will now illustrate how divergence-based cost works using the mountain car problem.

6.1. The Mountain Car Problem. Here, we use active inference and a generative model based on (31) to solve a fairly difficult problem in optimal control based purely on the theoretical treatment above. Crucially, the agent that solves this problem has no prior information about constraints on its action and yet it can respond adaptively, when supplied with a cost-function, to find its goal almost immediately. Note that there is no value-learning because we use divergence-based cost, which does not require a value-function. Furthermore, this behaviour is resilient to perturbations, because the policy provides predictions, which are fulfilled by action.

In the mountain car problem, one has to park a mountain car halfway up a hill. However, the car is not sufficiently powerful to ascend the hill directly. This means the only solution to problem is to back away from the parking location and then accelerate towards it, in the hope of acquiring sufficient momentum to access the target. The upper left panel of Figure 5 shows the landscape or potential energy function (with a minimum at position, $x = -0.5$) that exerts

forces on the car. The car is shown at the target position at the top of the hill at $x = 1$ (red dot). The equations of motion of the car are shown below. Crucially, at $x = 0$ the force is unity and cannot be overcome by the agent, because a squashing function $-1 \leq \sigma(a) \leq 1$ is applied to action. The right panels show the agent’s generative model in terms of the equations of motion in (13). These correspond to (31), where the cost-function is shown on the upper right. Here, the cost-function $c(x, z)$ has an auxiliary parameter $z \in \mathbb{R}$ that enables cost to be switched on or off. This can be thought of as satiety, such that when $z = 0$ cost is positive everywhere, except the target location (see figure legend for details). In this example, the sensory mapping and its assumed form were just $\mathbf{g}(\tilde{x}) = \tilde{x}$. Notice that the true equations of motion depend on action while the policy does not. This means the actual behaviour of the agent is selected from a family of flows by action. Figure 6 shows two exemplar flows under different values for action. Under active inference, action tries to realise conditional beliefs that are specified by the policy or empirical priors on motion.

Figure 7 shows how paradoxical but adaptive behaviour (moving away from a target to ensure it is secured later) emerges from this sort of generalised policy on the motion of hidden states, using $c(x, 0)$. These simulations are the results of integrating (1) and (2) (see Appendix H for details). The inferred hidden states (upper right) show that the car explores its landscape until it encounters the target and negative cost or friction increases dramatically to prevent it escaping (i.e., falling down the hill). This ensuing trajectory is shown on the upper left panel. The paler lines provide exemplar trajectories from other trials, with different starting positions. In the real world, friction is constant (one eighth). However, the car expects friction to change with position, enforcing exploration or exploitation. These expectations are fulfilled by action (lower right).

It is important to appreciate what has and what has not been achieved in this simulation: using a single scalar cost-function of position we have been able to induce adaptive goal-directed behaviour without any value-learning or enforced exploration of the environment. This behaviour is elicited immediately without the need for repeated trials or exposures. Furthermore, the requisite exploration and exploitation is manifest as a direct consequence of the agent’s priors, without the need for carefully calibrated stochastic terms during training. This sort of adaptive behaviour is reminiscent of foraging seen in insects and other biological agents; where the environment is explored under autonomous dynamics until some reward is encountered. In an experimental context, the above simulation could be regarded in terms of eliciting foot-shock escape behaviour [75]; in that the uniformly high cost can only be avoided by escaping to a low-cost location. It may seem implausible to specify behaviour in terms of cost that is generally high everywhere; however, one can see easily how drive states might correspond to high cost and subsequent foraging or seeking behaviour [76, 77]. Mathematically, this reflects the fact that cost plays a permissive role, in that it ensures maxima of the ergodic density lie in low-cost regions of state-space by precluding maxima elsewhere. In this sense, the emphasis is

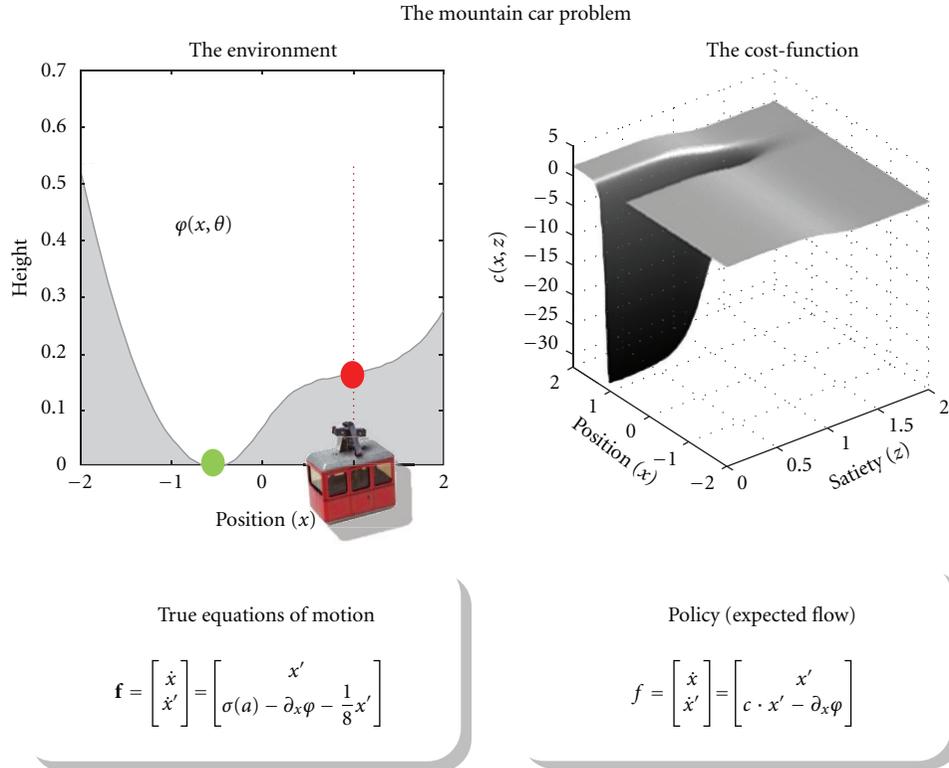


FIGURE 5: The mountain car problem. The upper left panel shows the landscape or potential energy function $\varphi(x, \theta)$, with a minimum at position, $x = -0.5$ (green dot) that exerts forces on the car. The car is shown at the target position at the top of the hill at $x = 1$ (red dot). The equations of motion of the car are shown below. Crucially, at $x = 0$ the force is unity and cannot be overcome by the agent, because a squashing function $-1 \leq \sigma(a) \leq 1$ is applied to action. This means the agent can only access the target by starting on the left hill to gain enough momentum to carry it up the other side. The right panels show the cost function and empirical priors (model of flow) that constitute the agent. Cost is a function of position and a hidden (e.g., physiological) state that plays a role of satiety $c(x, z) = (16 \cdot \exp(-64(x-1)^2) - 1) \cdot (\tanh(8(z-1)) - 1) - 1$. When satiety is high, cost is uniformly negative; $c(x, \infty) = -1$. Conversely, when satiety is low cost becomes negative near, and only near, the target location; $c(x, 0) = 1 - 32 \cdot \exp(-64(x-1)^2)$. The equations of motion on the lower right are constructed to ensure that fixed points are only stable in regions of negative cost or divergence: see main text.

on *autovitiating* of costly fixed points, as opposed to the *autopoiesis* of attractors.

What we have not addressed here is learning: we have assumed that the agent has already learned the potential energy function $\varphi(x)$ that corresponds to environmental constraints on its motion: see [6] for an example of this learning. However, there is a subtle but important point about learning: learning the parameters of the potential function corresponds to learning divergence-free flow, which does not affect the ergodic density or fixed points of the attractor. This contrasts with value-learning, in which divergence-flow is unspecified and the parameters of the value-function are learned. We now look at generalising divergence-based schemes and their role in prescribing sequential and itinerant dynamics.

6.2. Optimal Itinerancy and Weakly Attracting Sets. There are clearly many different ways in which we could formulate generalised policies and constrain them with cost-functions: We will concentrate on the use of cost to specify itinerant policies: itinerancy is important because it provides a principled explanation for exploration and foraging in ethology

[78]. Furthermore, it provides a key connection to dynamical systems theory approaches to the brain [79] that emphasise the importance of itinerant chaos [80], metastability [81], self-organised critically [82], winnerless competition [83], and attractors [84]. Similar constructs based on metastability have been invoked to understand the dynamics of molecular biology and the emergence of disease states like cancer [85]. The common theme here is the induction of itinerancy through the destruction of fixed-points or the gradients causing them [86]. The ensuing *attractor ruins* or relics [87] provide a framework for heteroclinic orbits that are ubiquitous in neurobiology, in electrophysiology [88], cognition [89], and large-scale neuronal dynamics [90].

It is fairly easy to extend the mountain car example above to produce itinerant behaviour with heteroclinic orbits and winnerless competition [83]. Intuitively, this can be regarded as adaptive behaviour in which various rewards are accessed in sequence to maintain physiological homeostasis (e.g., eating and drinking). This is straightforward to model by making the cost-function state dependent. This enables cost to be changed by the behaviours it induces. A simple example is provided in Figure 8, in which we have made the satiety

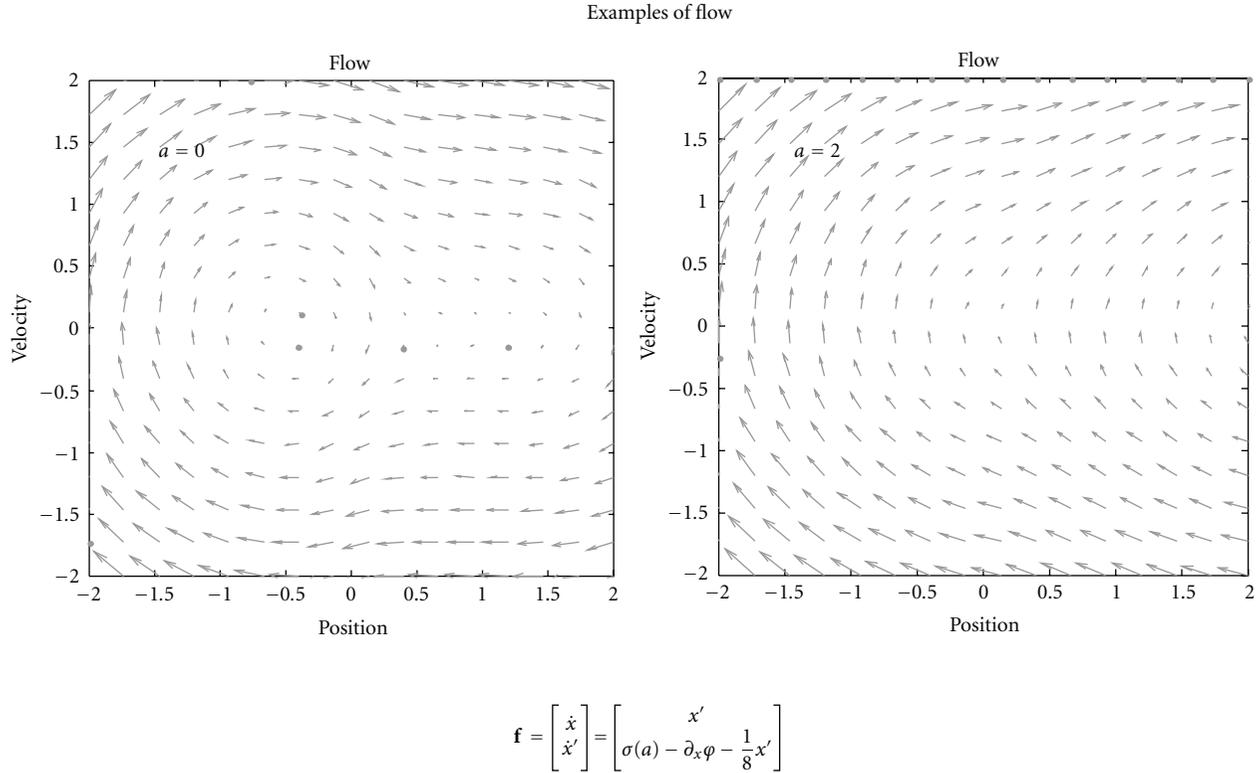


FIGURE 6: Examples of flow. This figure provides two examples of flow under two levels of action based on the equations of motion in the previous figure (the mountain car problem). These action-dependent flows provide a repertoire from which the agent has to compose a policy that conforms to its prior beliefs.

parameter of the cost function a hidden state in the generative model (and environment) so that satiety increases whenever cost is low. Cost can only be low in attractive states, which means attractive states become unattractive when occupied for too long. In the mountain car setup, when satiety rises, cost is uniformly low everywhere and the agent will simply settle at the bottom of the valley and stay there until satiety decays sufficiently to make the parking location attractive again. Figure 8 shows the equations of motion and ensuing dynamics, using the same format as in previous figures. This behaviour is characteristic of winnerless competition, in the sense that attractive fixed points are inherently unstable and release the trajectory to the next fixed point in the sequence. In this instance, instability is induced dynamically through state-dependent cost. This causes the mountain car to periodically rouse itself from the bottom of the valley and visit the parking location for a short time, until sated and then return to the bottom of the valley for a rest.

The vitiation of costly attractors is a mechanism that appears in several guises and has found important applications in a number of domains. For example, it is closely related to the notion of autopoiesis and self-organisation in situated (embodied) cognition [42]. It is formally related to the destruction of gradients in synergetic treatments of intentionality [86]. Mathematically, it finds a powerful application in universal optimisation schemes [91] and, indeed, as a model of perceptual categorization [92]. The dynamical phenomena, upon which these schemes rest, involve an itinerant

wandering through state-space along heteroclinic channels (orbits connecting different fixed points). Crucially, these attracting sets are weak (Milnor) attractors or attractor ruins that expel the state until it finds the next weak attractor. The result is a sequence of transitions through state-space that, in some instances, can be stable and repeating. The resulting stable heteroclinic channels have been proposed as a metaphor for neuronal dynamics and underlying cognitive processing [83]. Furthermore, the notion of Milnor or ruined attractors underlies much of the technical and cognitive literature on itinerant dynamics. For example, one can explain “a range of phenomena in biological vision, such as mental rotation, visual search, and the presence of multiple time scales in adaptation” using the concept of weakly attracting sets [92], see also [93]. It is this sort of dynamical behaviour that may underpin generalised policies that are specified directly in terms of equations of motion (as opposed to value functions in optimal control).

6.3. Summary. In this section, we have seen how cost can be used to induce attractive fixed points in hidden state-space while destroying unattractive fixed points. This does not involve any value-learning but rests upon the fact that stable fixed points result from flow with the negative divergence. We have seen how these policies can be realised in a straightforward manner under active inference and how endowing cost with a context sensitivity leads to itinerant but purposeful behaviour that is reminiscent of biological

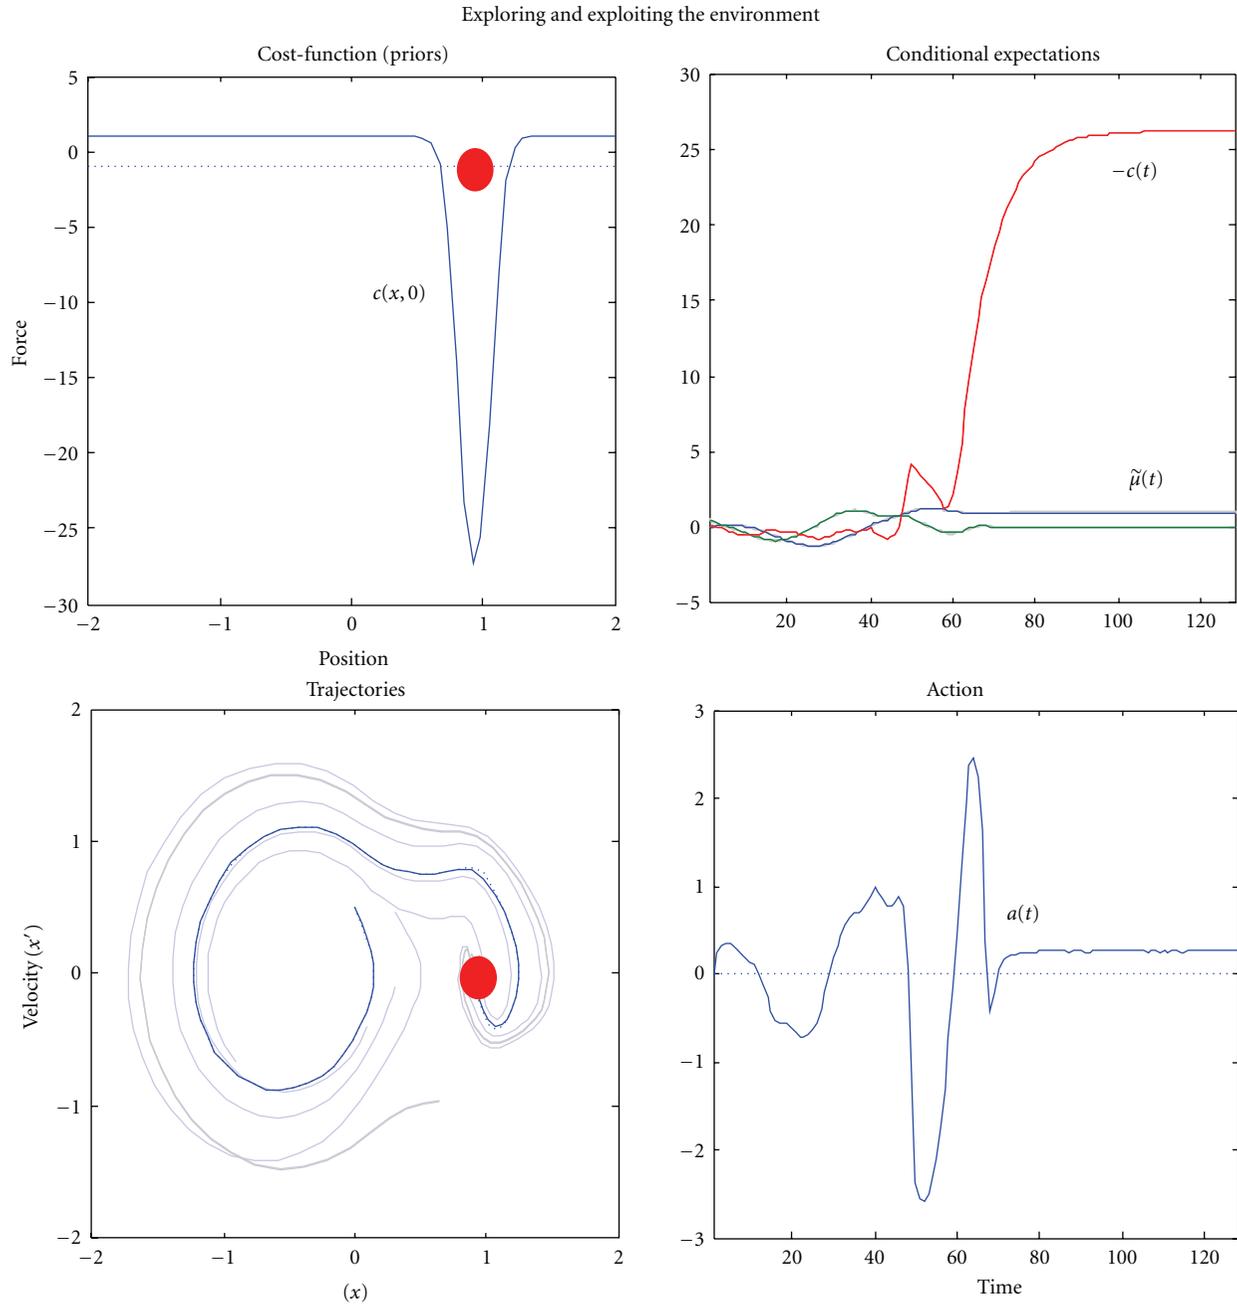


FIGURE 7: Active inference with generalised policies. This example shows how paradoxical but adaptive behaviour (moving away from a target to secure it later) emerges from simple priors on the motion of hidden states. These priors are encoded in a cost function $c(x, 0)$ (upper left). The form of the agent's (generalised) policy ensures that divergence is positive or friction is negative in regions of positive cost, such that the car expects to go *faster*. The inferred hidden states (upper right: position in blue, velocity in green, and friction in red) show that the car explores its landscape until it encounters the target and friction increases dramatically to prevent it escaping (i.e., falling down the hill). The ensuing trajectory is shown in blue (lower left). The paler lines provide exemplar trajectories from other trials, with different starting positions. In the real world, friction is constant (one eighth). However, the car expects friction to change with position, enforcing exploration or exploitation. These expectations are fulfilled by action (lower right), which tries to minimise free energy.

systems. The basic message here is that it may be sufficient to understand adaptive self-organised behaviour purely in terms of the itinerant dynamics induced by an agent's (implicit) prior beliefs about its motion through state-space. These dynamics and their associated attractors can be characterised in terms of unstable fixed points (weak attractors)

and, in most instances, an associated sequence of heteroclinic orbits. In this dynamical setting, a natural way to specify (and inherit) the weakly attracting sets that define phenotypic behaviour is to destroy or preclude (stable) fixed points that do not belong to attracting set. Note that this stands in stark contrast to optimal control theory, which tries

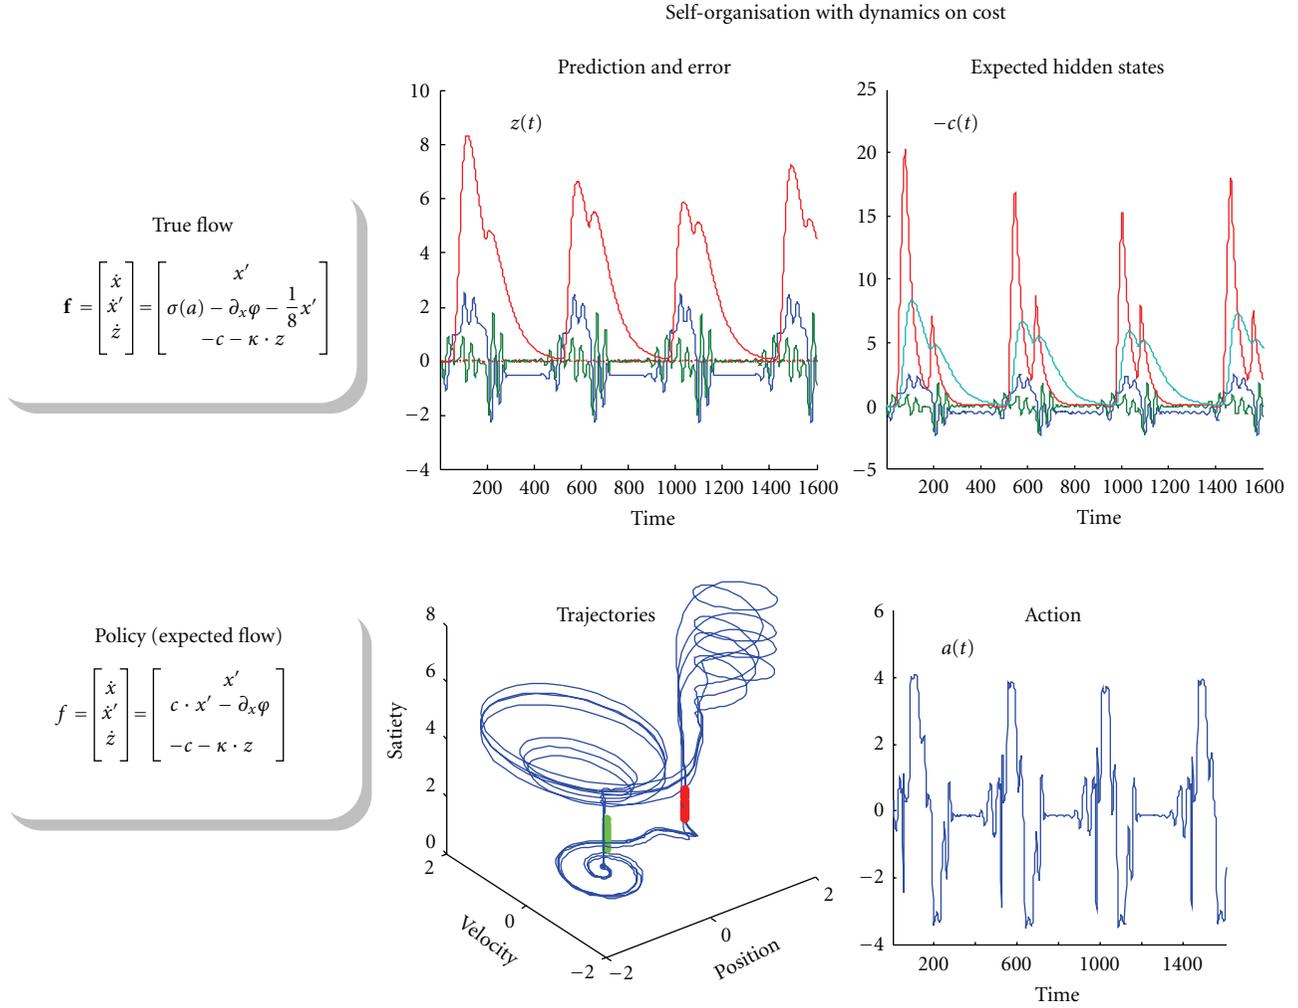


FIGURE 8: Optimal itinerancy. This figure shows how itinerant dynamics can be constrained by a cost function, leading to a stable heteroclinic channel, in which unstable but attractive fixed points are visited in succession. Here, we have exploited the specification of cost in terms of satiety, which has been made a hidden (physiological) state. This makes cost time dependent and sensitive to the recent history of the agent’s states. Placing dynamics on cost enables us to model sequential behaviour elicited by cost functions that are suppressed by the behaviour they elicit. The left panels show the true (upper) and modelled (lower) equations of motion on hidden states, where the latter are constrained by the cost function in Figure 5. Here, satiety increases with rewards (negative cost) and decays with first-order kinetics. The resulting behaviour is summarised in the right-hand side panels. The upper left panel shows the predictions of hidden states and prediction errors; where predictions are based upon the conditional beliefs about hidden states shown on the upper right. These predictions prescribe optimal action (lower right), which leads to the behavioural orbits shown on the lower left. The characteristic feature of the ensuing dynamics is a sequential return to unstable fixed points; denoted by the minimum of the potential landscape (green dots) and the cost-dependent (unstable) fixed point at the target location (red dots).

to optimise the flow using value-functions. However, the (generally intractable) computation of these functions may be unnecessary and unnatural, if it is sufficient to place straightforward constraints on the flow that defines value.

7. Discussion

This paper started with the existence of agents (systems) with defining characteristics that are conserved over time. We used arguments from ergodic theory, random dynamical systems and information theory to identify the imperatives for their dynamics. The picture that emerges can be summarised as

follows. Agents are equipped (by evolution or engineering) with a characteristic (cost) function of the states they should occupy or possess. This function places constraints on prior beliefs about motion through state-space (state transitions). Action realises this policy by suppressing random or surprising deviations from the ensuing predictions, thereby minimising surprise and the entropy of the ergodic density (over long periods of time). The result is a random dynamical attractor, with small measure that ensures agents occupy a limited attracting set of states (or expresses phenotypic traits that are conserved over time). Every policy (flow) has an associated value-function, which is the (log) ergodic density

(or negative surprise) of any generalised state. The self-consistency of cost-functions and the ergodic densities they engender is assured by natural selection; in the sense that cost-functions that do not induce ergodic behaviour cannot be inherited. In the final sections, we compared and contrasted policies from optimal control theory with generalised policies based on dynamical systems theory that lead to itinerant behaviour.

7.1. Dynamics versus Reinforcement-Learning. The formulations in this paper emphasise the link between cost-functions and policies. Optimal control theory and reinforcement-learning assumes value is expected cost in the future. This enables the policy to be optimised in terms of control, such that the expected path integral of cost is minimised. The ensuing policies are prescribed directly by value, which acts as a guiding function. This entails finding a solution to a set of self-consistent equations linking value and cost. However, this is a difficult problem and leads to some inconsistencies; for example, the autonomous or random explorations of state-space needed to furnish solutions to the Bellman equations are precisely the fluctuations that optimal control is trying to avoid. Generalised policies resolve these difficulties because they do not define value as expected cost: value is defined in terms of the states that are visited most frequently (i.e., the ergodic density), and is a function of flow (the policy). The last section tried to show that there are straightforward ways to place constrain policies; namely, to destroy unattractive fixed points. In summary, reinforcement-learning starts with a cost-function from which the value-function is derived. The value is then used to optimise a policy. Conversely, in the setting of random attractors, cost-functions constrain the policy directly. By definition, the policy then maximises value or minimises surprise. This eschews the solution of the appropriate Bellman equation, provides a principled explanation for exploratory or itinerant dynamics, and affords an efficient and biologically plausible scheme. Furthermore, it allows action to actively suppress unpredicted deviations from the policy.

The importance of dynamical itinerancy has been articulated many times in the past [94], particularly from the perspective of computation and autonomy; see [93] for a focus on Milnor attractors. It has also been considered formally in relation to cognition; see [87] for a focus on attractor relics, ghosts, or ruins. Indeed, there is growing interest in understanding brain dynamics *per se* in terms of itinerancy and metastability [81, 83, 88, 89]. Tani et al., [95] consider itinerant dynamics in terms of bifurcation parameters that generate multiple goal-directed actions (on the behavioural side) and optimization of the same parameters (when recognizing actions). They provide a series of elegant robotic simulations to show generalization by learning with this scheme. See also [96] for interesting simulations of itinerant exploration, using just prediction errors on sensory samples over time.

Reinforcement-learning frames the problem of adaptive behaviour in terms of accessing distal and sparse rewards. In one sense this is not a problem; it is the solution entailed by an agent and its interactions with the environment. In this view, agents do not seek out valuable (rewarding) states;

valuable states are just states the agent frequents. This challenges any preconception that optimal control has a central or unique role in adaptive behaviour. Having said this, the premise of optimal control and reinforcement-learning that agents minimise expected future costs is a compelling and enduring heuristic. This heuristic may be exploited by the brain, particularly in terms of high-level (e.g., cognitive) processing using model-based schemes.

7.2. Value-Learning versus Perceptual Learning. The mountain car example can be regarded as a model of behavioural responses constrained by *innate or formal* priors (cost-functions). However, most of the interesting issues in a biological setting may rest on *acquired or empirical* priors that are optimised during perception. Irrespective of the details of this optimisation, under active inference, optimisation becomes a perceptual inference and learning problem. In other words, the notion that stimulus-response links are selectively reinforced during learning disappears and is replaced by the learning of stimulus-stimulus associations. These prescribe conditional beliefs, which action fulfills. This is important because it places much of behaviourism in the domain of perception and reiterates the close links between action and perception.

It is well known from the complete class theorem that there is a close relationship between priors and cost-functions; in the sense that any admissible decision rule is Bayes-optimal for at least one prior and cost-function [97]. The treatment in this paper suggests that when decisions involve inferred states of the world, cost-functions can be treated as priors. Heuristically, cost-functions are a fixed attribute of the agent and can therefore only manifest as formal priors on the agent's inference and consequent behaviour. This is particularly important in a biological setting, where natural selection endows agents with formal or innate priors that constrain their exchanges with the environment.

8. Conclusion

In this paper, we have tried to understand optimal control theory in relation to the free energy principle. We started with a review of ensemble dynamics and the perspective it provides on reinforcement-learning and optimal control. These approaches furnish policies or equations of motion that converge on attractors in state-space that are specified by a cost-function. Conventional schemes specify these equations in terms of value-functions or cost-to-go, which entail the solution of the appropriate Bellman equation. We considered a dynamical alternative based on the selective destruction of stable fixed-points in costly regimes of state-space. Although less efficient at minimising the path integral of cost, this divergence-based scheme involves no value-learning and accounts for exploratory dynamics that do not need stochastic interventions. In this context, the policies of optimal control become formal priors in generative models used to infer hidden states and predict sensations. Action fulfils these predictions by suppressing a free energy bound on surprise. Crucially, optimising action, perceptual inference, perceptual learning, and the priors themselves are all mandated by the

free energy principle. This principle is simply a restatement of the fact that adaptive systems resist a natural tendency to disorder. In summary, agents must be equipped with policies that prescribe their expected motion through state-space (i.e., state transitions) and the ability to counter surprising (random) deviations from the expected trajectory through action. The motion prescribed by the policy (and realised by action) induces low entropy densities (in terms of ensemble dynamics) or random attractors with small measure (in terms of random dynamical systems). These are sufficient to explain the existence of ergodic self-organising systems, whose attributes are conserved over time.

Appendices

A. Entropy Production

This appendix shows that the dispersion of states $x \in X$ by random fluctuation increases entropy. This implies (autopoietic) flow must reduce entropy to maintain an invariant probability density over states. This is based on the following lemma.

Lemma A.1 (entropy production). *The entropy production of a differentiable dynamical system (T, X, f) can be decomposed into a flow and non-negative dispersion term*

$$\dot{\mathcal{H}}(X) = \int_X \nabla \cdot (fp) \ln p(x) dx + \int_X \frac{\nabla p \cdot \Gamma \cdot \nabla p}{p(x)} dx. \quad (\text{A.1})$$

Proof. We start with the Fokker-Planck equation describing the rate of change of an ensemble density $p(x, t)$ in terms of dispersion due to random fluctuations and flow $f : X \rightarrow X$,

$$\dot{p}(x, t) = \nabla \cdot \Gamma \nabla p - \nabla \cdot (fp). \quad (\text{A.2})$$

Without loss of generality, we assume some diffeomorphism of state-space in which the diffusion tensor can be expressed as $\Gamma(x) = \gamma(x)I$, where $\gamma(x) \in \mathbb{R}^+$. The existence of this state-space is assured by the positive definiteness of the diffusion tensor. The entropy is

$$\begin{aligned} \mathcal{H}(X) &= - \int p \ln p dx \implies \\ \dot{\mathcal{H}}(X) &= - \int \dot{p} \ln p dx - \int \dot{p} dx = - \int \dot{p} \ln p dx. \end{aligned} \quad (\text{A.3})$$

The second term in the expression for entropy production $\dot{\mathcal{H}}(X)$ disappears because probability mass is conserved. The rate of change of entropy can be decomposed into flow- and dispersion-dependent terms by substituting (A.2) into (A.3)

$$\dot{\mathcal{H}}(X) = \int \nabla \cdot (fp) \ln p dx - \gamma \int (\nabla^2 p) \ln p dx. \quad (\text{A.4})$$

Finally, the self-adjoint property of the Laplacian operator ∇^2 means that the second (dispersion) term is greater or

equal to zero:

$$\begin{aligned} \int (\nabla^2 p) \ln p dx &= \int p (\nabla^2 \ln p) dx \\ &= \int (\nabla^2 p) dx - \int \frac{\nabla p \cdot \nabla p}{p} dx \quad (\text{A.5}) \\ &= - \int \frac{\nabla p \cdot \nabla p}{p} dx \leq 0. \end{aligned}$$

The integral of the curvature $\nabla^2 p$ disappears because $p := p(x, t)$ is a proper density. Substituting (A.5) into (A.4) gives (A.1). \square

Remarks A.2. Equation (A.5) means that random fluctuations increase entropy in proportion to their amplitude and the roughness of the ensemble density. In the absence of deterministic flow, the entropy increases until the density has dispersed and its gradients disappear; that is, $\nabla p = 0 \Rightarrow \dot{\mathcal{H}}(X) = 0$. In the presence of flow, (A.4) implies, at steady state $\dot{\mathcal{H}}(X) = 0$ and flow decreases entropy to offset dispersion-related increases. See Tomé [41] for a fuller treatment in terms of probability currents.

B. Sensory Entropy

This appendix shows that the entropy of hidden states is bounded by the entropy of sensory states. This means that if the entropy of generalised sensory signals is minimised, so is the entropy of the hidden states that caused them. We will assume sensory states are an analytic function of hidden states plus some fluctuations. This implies (in generalised coordinates of motion) that

$$\begin{aligned} \begin{matrix} s \\ s' \\ s'' \\ \vdots \end{matrix} &= \begin{matrix} g(x) \\ \partial_x g \cdot x' \\ \partial_x g \cdot x'' \\ \vdots \end{matrix} + \begin{matrix} \omega_s \\ \omega'_s \\ \omega''_s \\ \vdots \end{matrix}, \end{aligned} \quad (\text{B.1})$$

where the second equality can be written more compactly as $\tilde{s} = \tilde{g}(\tilde{x}) + \tilde{\omega}$.

Lemma B.1 (entropy bounds). *If the sensitivity of the sensory mapping $|\partial \tilde{g} / \partial \tilde{x}|$ is uniform over hidden states, then the entropy of hidden states is bounded by the entropy of sensory states, to within an additive constant:*

$$H(S | m) \geq H(X | m) + \ln \left| \frac{\partial \tilde{g}}{\partial \tilde{x}} \right|. \quad (\text{B.2})$$

Proof. Because the random fluctuations are conditionally independent, we have (Theorem 6.5 in [98, page 151]) the following:

$$H(\tilde{S} | m) \geq H(\tilde{X} | m) + \int p(\tilde{x} | m) \ln \left| \frac{\partial \tilde{g}}{\partial \tilde{x}} \right| d\tilde{x}. \quad (\text{B.3})$$

Because the log sensitivity is constant, it can be placed outside the integral, which integrates to unit, giving (B.2). \square

Remarks B.2. Here, $\partial\tilde{g}/\partial\tilde{x}$ is the sensitivity or gradient of the generalised sensory mapping with respect to the hidden states. The integral in (B.2) reflects the fact that entropy is not invariant to a change of variables and assumes that the sensory mapping $\tilde{g}: X \rightarrow S$ is diffeomorphic (i.e., bijective and smooth). This requires the hidden and sensory state-spaces to have the same dimension, which can be assured by truncating generalised states at an appropriate order to make $\dim(S | m) = \dim(X | m)$.

C. The Laplace Assumption

This appendix shows why it is only necessary to retain the conditional mean under a fixed-form Laplace assumption about the recognition density. If we assume $q(\vartheta) = \mathcal{N}(\tilde{\mu}, \Sigma)$ is Gaussian (the Laplace assumption), then we can express free energy in terms of its sufficient statistics; the mean and covariance, $(\tilde{\mu}, \Sigma)$ using Gibb's energy $U(t) = -\ln p(\tilde{s}, \vartheta | m)$ and $n = \dim(\tilde{\mu})$

$$\mathcal{F} = U(\tilde{\mu}) + \frac{1}{2} \text{tr} \left(\Sigma \frac{\partial}{\partial \tilde{\mu}^2} U \right) - \frac{1}{2} \ln |\Sigma| - \frac{n}{2} \ln 2\pi e. \quad (\text{C.1})$$

We can now minimise free-energy with respect to the conditional precision $\Pi = \Sigma^{-1}$ by solving $\partial_{\Sigma} \mathcal{F} = 0$ to give

$$\partial_{\Sigma} \mathcal{F} = -\frac{1}{2} \Pi + \frac{1}{2} \frac{\partial}{\partial \tilde{\mu}^2} U = 0 \implies \Pi = \frac{\partial^2}{\partial \tilde{\mu}^2} U(\tilde{\mu}). \quad (\text{C.2})$$

Critically, this is an analytic function of the mean and does not have to be evaluated explicitly. Furthermore, we can simplify the expression for free-energy by eliminating the curvature from (C.1) to give $\mathcal{F} = U(\mu) - (1/2) \ln |\Sigma| - (n/2) \ln 2\pi$.

D. Value and Ergodic Densities

For simplicity, we will ignore the state-dependence of $Q(x)$ and assume it is constant; a full treatment can be found in [58].

Lemma D.1 (ergodic density). *The ergodic density of a dynamical system with Fokker Planck operator $\Lambda = \nabla \cdot (\Gamma \nabla - f)$ and flow $f = (\Gamma + Q) \nabla V$ subject to $Q = -Q^T$ is given by*

$$p(x | m) = \exp(V) \implies \nabla p = p \nabla V \implies \Lambda p = 0. \quad (\text{D.1})$$

Proof. By substituting $\nabla p = p \nabla V$ and $f = \Gamma \nabla V + Q \nabla V$ into the Fokker Planck operator we get

$$\begin{aligned} \Lambda p &= \nabla \cdot \Gamma \nabla p - \nabla \cdot (f p) \\ &= \nabla \cdot (p \Gamma \nabla V) - \nabla \cdot (p \Gamma \nabla V) - \nabla \cdot (p Q \nabla V) \\ &= -p \nabla \cdot (Q \nabla V) - (Q \nabla V) \cdot \nabla p \\ &= -p (\nabla \cdot (Q \nabla V) + (Q \nabla V) \cdot \nabla V) \end{aligned} \quad (\text{D.2})$$

one can see that (D.2) is satisfied when the second component of flow $Q \nabla V$ is divergence free and orthogonal to ∇V .

Given $Q = -Q^T$, it is easy to see both these conditions are met

$$\begin{aligned} \nabla \cdot (Q \nabla V) &= \text{tr}(Q \partial_{xx} V) = 0, \\ \nabla V \cdot (Q \nabla V) &= \text{tr}(Q \partial_x V \partial_x V^T) = 0. \end{aligned} \quad (\text{D.3})$$

This means that $p = \exp(V)$ is the ergodic density or eigen-solution $\Lambda p = 0$ of the Fokker-Planck operator describing density dynamics. \square

E. Cost Functions and Value

Here, we derive cost as a function of value, using discrete and continuous time formulations.

Discrete Case. For the discrete case, let \mathbf{c} (resp., \mathbf{v}) be a row-vector of the cost (resp., value) of every discrete state $x_t \in X$ at time t and $\mathbf{P}_{ij} = p(x_{t+1} = i | x_t = j)$ be the transition probability or policy. By definition, value is (assuming $\mathbf{cP}^\infty = 0$)

$$\begin{aligned} \mathbf{v} &= -\sum_{t=0}^{\infty} \mathbf{cP}^t \implies \\ \mathbf{vP} &= -\sum_{t=1}^{\infty} \mathbf{cP}^t = \mathbf{v} + \mathbf{c} \implies \mathbf{c} = \mathbf{vP} - \mathbf{v} = \mathbf{v}_1 - \mathbf{v}_0. \end{aligned} \quad (\text{E.1})$$

The subscripted $\mathbf{v}_t := \mathbf{vP}^t$ is the value expected at time step t in the future. It can be seen that cost is just the expected increase in value over the next time step.

Continuous Case. In the continuous time case, the value of $x_0 \in X$ is the cost expected over future times $t \in (t_0, \infty]$, starting with a point density, $p(x, t_0 | m) = \delta(x_0)$. In the infinite horizon case, assuming the cost expected at equilibrium is zero; that is,

$$\int_{x \in X} c(x) p(x | m) dx = 0. \quad (\text{E.2})$$

We have, by definition

$$\begin{aligned} V(x_0) &= -\int_{t_0}^{\infty} \int_{x \in X} c(x) p(x, t | m) dx dt \implies \\ \dot{V}(x_0) &= \int c(x) p(x, t_0 | m) dx = c(x_0) \\ &= \int V(x) \dot{p}(x, t_0 | m) dx \\ &= \int V(x) (\nabla \cdot \Gamma \cdot \nabla \delta(x_0) - f \cdot \nabla \delta(x_0) \\ &\quad - \delta(x_0) \nabla \cdot f) dx \\ &= \nabla \cdot \Gamma \cdot \nabla V(x_0) + \nabla \cdot (V(x_0) f) \\ &\quad - V(x_0) \nabla \cdot f \implies \\ c(x) &= f \cdot \nabla V(x) + \nabla \cdot \Gamma \cdot \nabla V(x) = \dot{V}(x). \end{aligned} \quad (\text{E.3})$$

This says that the increase in value expected over time is the cost at any location in state-space. However, this is also the increase in expected value over state-space, which can

be expressed using the Fokker-Planck equation (A.2). Evaluating this expression at the initial density shows that cost comprises two components; one due to deterministic flow and another to dispersion (the first and second terms in the last equality).

F. Optimal Control and Policies

This appendix provides a brief review of classical formulations of optimal control in discrete and continuous time. For simplicity, we will ignore random fluctuations in the continuous case. The value-function prescribes an optimal policy, which can be expressed in discrete and continuous form as

$$\begin{aligned}\pi^* &= \arg \max_{\pi} \mathbf{vP}(\pi), \\ \pi^* &= \arg \max_{\pi} \mathbf{f}(x, \pi) \cdot \nabla V(x),\end{aligned}\tag{F.1}$$

where π parameterises the transition probability matrix in the discrete case and $u = \pi(x)$ represents a mapping between states and control in the continuous case. Equation (F.1) means that the optimal policy $\mathbf{P}(\pi^*)$ maximises the expected value of the next state. This is the solution to the Bellman and Hamilton-Jacobi-Bellman equations for the discrete and continuous cases, respectively,

$$\begin{aligned}\mathbf{c} &= \max_{\pi} \{\mathbf{vP}(\pi) - \mathbf{v}\}, \\ c(x) &= \max_{\pi} \{\mathbf{f}(x, \pi) \cdot \nabla V(x)\}.\end{aligned}\tag{F.2}$$

More generic forms of the Bellman and HJB equations consider loss incurred by the state-transitions *per se*, in which case the cost-function is brought into the maximum operator

$$\begin{aligned}\max_{\pi} \{\mathbf{vP}(\pi) - \mathbf{c}(\pi) - \mathbf{v}\} &= 0, \\ \max_{\pi} \{\mathbf{f}(x, \pi) \cdot \nabla V - c(x, \pi)\} &= 0.\end{aligned}\tag{F.3}$$

Furthermore, one might want to consider cost-functions that change with time giving, for the continuous case

$$\max_{\pi} \{\mathbf{f}(x, \pi) \cdot \nabla V - c(x, \pi)\} = -\dot{V}.\tag{F.4}$$

Although important cases in engineering problems, these generalisations are not so interesting from the point of view of ensemble dynamics with ergodicity (especially when one considers cost-functions on generalised states). However, in differential game theory, the HJB equation generalizes to the Hamilton-Jacobi-Isaacs (HJI) equations [99].

G. Value Learning and Optimality Equations

This appendix reviews standard solutions to the Bellman (discrete) and HJB (continuous) equations based on the Robbins-Monro algorithm (or stochastic iteration algorithm [67]) for solving systems of the form $\mathbf{v} = h(\mathbf{v})$,

$$\mathbf{v} \leftarrow \mathbf{v} + \rho \delta : \delta = h(\mathbf{v}) - \mathbf{v}.\tag{G.1}$$

Here, δ corresponds to a prediction error and $\rho \in [0, 1]$ determines the convergence rate.

Discrete Case. Using $h(\mathbf{v}) = \mathbf{vP}(\pi^*) - \mathbf{c}$ from (F.3), one could optimize the value-function of all states simultaneously using (G.1) [100]. This is an example of a *model-based* scheme because it rests on a model of state-transitions implicit in the policy $\mathbf{P}(\pi)$. Simpler *model-free* schemes adopt a stochastic or sampling approach to approximating the value-function with an estimate $v(x_t)$ of each state visited

$$v(x_t) \leftarrow v(x_t) + \rho \delta_t : \delta_t = v(x_{t+1}) - v(x_t) - c(x_t).\tag{G.2}$$

Here, δ corresponds to a reward prediction error. When its average converges to zero, then the long-term average $\mathbf{v}_i = \langle v(x_i) \rangle_t$ is the desired solution to the Bellman equation,

$$\langle \delta_t \rangle_t = 0 \implies c(x_t) = \langle v(x_{t+1}) - v(x_t) \rangle_t \implies \mathbf{c} = \mathbf{vP}(\pi^*) - \mathbf{v}.\tag{G.3}$$

The variance of $v(x_i)$ can be made arbitrary small by decreasing ρ . Q-learning rests on a similar scheme, but replaces value with quality $Q : X \times U \rightarrow \mathbb{R}$ on the joint space of states and action

$$\begin{aligned}Q(x_t, u_t) &\leftarrow Q(x_t, u_t) + \rho \delta : \delta = \max_u Q(x_{t+1}, u) - Q(x_t, u_t) \\ &\quad - c(x_t, u_t).\end{aligned}\tag{G.4}$$

From this stochastic value iteration scheme, one can calculate expected reward for any state-action pair. A related state-action-reward-state-action scheme [101] is

$$\begin{aligned}Q(x_t, u_t) &\leftarrow Q(x_t, u_t) + \rho \delta : \delta = Q(x_{t+1}, u_{t+1}) - Q(x_t, u_t) \\ &\quad - c(x_t, u_t).\end{aligned}\tag{G.5}$$

This learns the Q-values associated with taking the policy it follows while Q-learning learns the Q-values associated with the optimal policy. While these are important generalisations, one can simplify things theoretically by replacing Q-values with a value-function on the Cartesian product of state and action spaces; that is, $Q(x, u) \equiv V(x) : x \in X \times U$. This implicit factorisation of state-space has been used for approximating value-functions in Markov decision processes with large state and action spaces [102].

Continuous Case. A neurobiologically plausible continuous-time scheme was introduced in [72]. This scheme exploits the fact that the derivative of the value-functional (time-varying value of the current state) is just cost $\dot{V}(x(t)) = c(x)$ (in the deterministic limit of (E.3)). Here, the prediction-error or teaching signal $\delta = \dot{v}_c + \dot{v}_a$ is the derivative of *innate* and *acquired* value (e.g., the phasic responses of dopaminergic or cholinergic neurons). Innate value is the antiderivative of reward $\dot{v}_c = -c(x)$ registered by specific neuronal systems (e.g., the lateral hypothalamic area). Acquired value v_a is parameterised by connection strengths coupling sensory states to the neuronal structures that encode it (e.g., the amygdala). Connection strengths mediating acquired value and stimulus-response links are optimized by conventional associative plasticity. Crucially, this plasticity is modulated or enabled by the teaching signal. After convergence,

$$\delta = \dot{v}_c + \dot{v}_a = 0 \implies \dot{v}_a = -\dot{v}_c = c(x) = \dot{V}(x(t)),\tag{G.6}$$

the acquired value becomes the true value $\dot{v}_a = \dot{V}(x(t))$ (to within an additive constant).

H. Integrating Active Inference Schemes

The simulations in this paper involve integrating time-varying states in both the environment and the agent as a single system, which can be modelled with the following ordinary differential equation, where random fluctuations enter as analytic forcing terms and defining $\tilde{\mathbf{f}} := \mathbf{f}(\tilde{x}, a, \theta)$:

$$\dot{\mathbf{u}} = \begin{bmatrix} \dot{\tilde{s}} \\ \dot{\tilde{x}} \\ \dot{\tilde{\omega}}_s \\ \dot{\tilde{\omega}}_x \\ \dot{\tilde{\mu}} \\ \dot{a} \end{bmatrix} = \begin{bmatrix} \mathcal{D}\tilde{\mathbf{g}} + \mathcal{D}\tilde{\omega}_s \\ \tilde{\mathbf{f}} + \tilde{\omega}_x \\ \mathcal{D}\tilde{\omega}_s \\ \mathcal{D}\tilde{\omega}_x \\ \mathcal{D}\tilde{\mu} - \partial_{\tilde{\mu}}\mathcal{F} \\ -\partial_a\mathcal{F} \end{bmatrix}. \quad (\text{H.1})$$

To update these collective states we use a local linearization; $\Delta \mathbf{u} = (\exp(\Delta t \mathcal{J}) - I)\mathcal{J}(t)^{-1}\dot{\mathbf{u}}$ over time steps of Δt , where

$$\mathcal{J} = \frac{\partial \dot{\mathbf{u}}}{\partial \mathbf{u}} = \begin{bmatrix} 0 & \mathcal{D}\partial_{\tilde{x}}\tilde{\mathbf{g}} & \mathcal{D} & & & \mathcal{D}\partial_a\tilde{\mathbf{g}} \\ & \partial_{\tilde{x}}\tilde{\mathbf{f}} & I & & & \partial_a\tilde{\mathbf{f}} \\ & & \mathcal{D} & & & \\ & & & \mathcal{D} & & \\ -\partial_{\tilde{\mu}s}\mathcal{F} & & & \mathcal{D} - \partial_{\tilde{\mu}\tilde{\mu}}\mathcal{F} & -\partial_{\tilde{\mu}a}\mathcal{F} & \\ -\partial_{a\tilde{s}}\mathcal{F} & & & -\partial_{a\tilde{\mu}}\mathcal{F} & -\partial_{aa}\mathcal{F} & \end{bmatrix}. \quad (\text{H.2})$$

Because action can only affect free energy through sensory data, it can only affect sensory prediction error. Therefore action dynamics are as follows:

$$\begin{aligned} \dot{a} &= -\partial_a\mathcal{F} = -\partial_a\tilde{s} \cdot \tilde{\varepsilon}_s, \\ \partial_a\tilde{s} &= \partial_{\tilde{x}}\tilde{\mathbf{g}} \sum_i \mathcal{D}^{-i} (\partial_{\tilde{x}}\tilde{\mathbf{f}})^{i-1} \partial_a\tilde{\mathbf{f}}. \end{aligned} \quad (\text{H.3})$$

The partial derivative of the generalised sensory data with respect to action depends on changes in the generalised motion of hidden states and is specified by the generative process. Note that this partial derivative depends on the unknown generative process. However, sensations are generally produced in a simple and direct way by action (e.g., proprioception with stretch receptors), such that the dependency can be treated as known.

These equations may look complicated but can be evaluated automatically using numerical derivatives. All the simulations in this paper used just one routine—**spm_ADEM.m**. Demonstrations of this scheme are available as part of the SPM software (<http://www.fil.ion.ucl.ac.uk/spm>; **DEM_demo.m**), which reproduces the example shown in Figure 7.

Funding

The Wellcome Trust funded this paper.

References

- [1] K. Friston, J. Kilner, and L. Harrison, “A free energy principle for the brain,” *Journal of Physiology Paris*, vol. 100, no. 1–3, pp. 70–87, 2006.
- [2] R. S. Sutton and A. G. Barto, “Toward a modern theory of adaptive networks: Expectation and prediction,” *Psychological Review*, vol. 88, no. 2, pp. 135–170, 1981.
- [3] N. D. Daw and K. Doya, “The computational neurobiology of learning and reward,” *Current Opinion in Neurobiology*, vol. 16, no. 2, pp. 199–204, 2006.
- [4] P. Dayan and N. D. Daw, “Decision theory, reinforcement learning, and the brain,” *Cognitive, Affective and Behavioral Neuroscience*, vol. 8, no. 4, pp. 429–453, 2008.
- [5] Y. Niv and G. Schoenbaum, “Dialogues on prediction errors,” *Trends in Cognitive Sciences*, vol. 12, no. 7, pp. 265–272, 2008.
- [6] K. J. Friston, J. Daunizeau, J. Kilner, and S. J. Kiebel, “Action and behavior: a free-energy formulation,” *Biological Cybernetics*, vol. 102, no. 3, pp. 227–260, 2010.
- [7] K. J. Friston, J. Daunizeau, and S. J. Kiebel, “Reinforcement learning or active inference?” *PLoS ONE*, vol. 4, no. 7, Article ID e6421, 2009.
- [8] J. Nash, “Equilibrium points in n-person games,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 36, pp. 48–49, 1950.
- [9] C. F. Camerer, “Behavioural studies of strategic thinking in games,” *Trends in Cognitive Sciences*, vol. 7, no. 5, pp. 225–231, 2003.
- [10] P. Dayan, G. E. Hinton, R. M. Neal, and R. S. Zemel, “The Helmholtz machine,” *Neural computation*, vol. 7, no. 5, pp. 889–904, 1995.
- [11] T. S. Lee and D. Mumford, “Hierarchical Bayesian inference in the visual cortex,” *Journal of the Optical Society of America A*, vol. 20, no. 7, pp. 1434–1448, 2003.
- [12] R. P. N. Rao and D. H. Ballard, “Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects,” *Nature Neuroscience*, vol. 2, no. 1, pp. 79–87, 1999.
- [13] K. Friston, “A theory of cortical responses,” *Philosophical Transactions of the Royal Society B*, vol. 360, no. 1456, pp. 815–836, 2005.
- [14] R. Bellman, “On the theory of dynamic programming,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 38, pp. 716–719, 1952.
- [15] L. Demetrius, “Thermodynamics and evolution,” *Journal of Theoretical Biology*, vol. 206, no. 1, pp. 1–16, 2000.
- [16] A. Traulsen, J. C. Claussen, and C. Hauert, “Coevolutionary dynamics in large, but finite populations,” *Physical Review E*, vol. 74, no. 1, Article ID 011901, 2006.
- [17] S. Conway Morris, *Life’s Solution: Inevitable Humans in a Lonely Universe*, Cambridge University Press, Cambridge, UK, 2005.
- [18] R. D. McKelvey and T. R. Palfrey, “Quantal response equilibria for normal form games,” *Games and Economic Behavior*, vol. 10, no. 1, pp. 6–38, 1995.
- [19] H. Crauel and F. Flandoli, “Attractors for random dynamical systems,” *Probability Theory and Related Fields*, vol. 100, no. 3, pp. 365–393, 1994.
- [20] G. D. Birkhoff, “Proof of the ergodic theorem,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 17, pp. 656–660, 1931.

- [21] V. L. Ginzburg and L. D. Landau, "On the theory of superconductivity," *Zhurnal Eksperimental'noi i Teoreticheskoi Fiziki*, vol. 20, p. 1064, 1950.
- [22] H. Haken, *Synergetics: An Introduction. Non-Equilibrium Phase Transition and Self-Selforganisation in Physics Chemistry and Biology*, Springer, Berlin, Germany, 3rd edition, 1983.
- [23] J. Carr, *Applications of Centre Manifold Theory*, Springer, Berlin, Germany, 1981.
- [24] S. De Monte, F. d'Ovidio, and E. Mosekilde, "Coherent regimes of globally coupled dynamical systems," *Physical Review Letters*, vol. 90, no. 5, pp. 054102/1–054102/4, 2003.
- [25] R. V. N. Melnik and A. H. Roberts, "Computational models for multi-scale coupled dynamic problems," *Future Generation Computer Systems*, vol. 20, no. 3, pp. 453–464, 2004.
- [26] M. J. Davis, "Low-dimensional manifolds in reaction-diffusion equations. 1. Fundamental aspects," *Journal of Physical Chemistry A*, vol. 110, no. 16, pp. 5235–5256, 2006.
- [27] G. Deco, V. K. Jirsa, P. A. Robinson, M. Breakspear, and K. Friston, "The dynamic brain: from spiking neurons to neural masses and cortical fields," *PLoS Computational Biology*, vol. 4, no. 8, Article ID e1000092, 2008.
- [28] A. Hu, Z. Xu, and L. Guo, "The existence of generalized synchronization of chaotic systems in complex networks," *Chaos*, vol. 20, no. 1, Article ID 014001CHA, 10 pages, 2010.
- [29] J. Ito and K. Kaneko, "Self-organized hierarchical structure in a plastic network of chaotic units," *Neural Networks*, vol. 13, no. 3, pp. 275–281, 2000.
- [30] W. R. Ashby, "Principles of the self-organizing dynamic system," *The Journal of General Psychology*, vol. 37, pp. 125–128, 1947.
- [31] G. Nicolis and I. Prigogine, *Self-Organization in Non-Equilibrium Systems*, John Wiley & Sons, New York, NY, USA, 1977.
- [32] C. C. Moore, "Ergodicity of flows on homogeneous spaces," *American Journal of Mathematics*, vol. 88, pp. 154–178, 1966.
- [33] T. D. Frank, *Nonlinear Fokker-Planck Equations: Fundamentals and Applications*, Springer Series in Synergetics, Springer, Berlin, Germany, 2004.
- [34] K. Friston, "Hierarchical models in the brain," *PLoS Computational Biology*, vol. 4, no. 11, Article ID e1000211, 2008.
- [35] H. Qian and D. A. Beard, "Thermodynamics of stoichiometric biochemical networks in living systems far from equilibrium," *Biophysical Chemistry*, vol. 114, no. 2-3, pp. 213–220, 2005.
- [36] H. Qian, "Entropy demystified: the "thermo"-dynamics of stochastically fluctuating systems," *Methods in Enzymology*, vol. 467, pp. 111–134, 2009.
- [37] C. Kwon, P. Ao, and D. J. Thouless, "Structure of stochastic dynamics near fixed points," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, no. 37, pp. 13029–13033, 2005.
- [38] P. Ao, "Global view of bionetwork dynamics: adaptive landscape," *Journal of Genetics and Genomics*, vol. 36, no. 2, pp. 63–73, 2009.
- [39] H. Crauel, "Global random attractors are uniquely determined by attracting deterministic compact sets," *Annali di Matematica Pura ed Applicata*, vol. 4, pp. 57–72, 1999.
- [40] D. J. Evans, "A non-equilibrium free energy theorem for deterministic systems," *Molecular Physics*, vol. 101, no. 10, pp. 1551–1554, 2003.
- [41] T. Tomé, "Entropy production in nonequilibrium systems described by a Fokker-Planck equation," *Brazilian Journal of Physics*, vol. 36, no. 4, pp. 1285–1289, 2006.
- [42] H. R. Maturana and F. Varela, "Autopoiesis: the organization of the living," in *Autopoiesis and Cognition*, F. Varela and H. R. Maturana, Eds., Reidel, Dordrecht, The Netherlands, 1980.
- [43] R. P. Feynman, *Statistical Mechanics*, Benjamin, Reading, Mass, USA, 1972.
- [44] G. E. Hinton and D. van Camp, "Keeping neural networks simple by minimizing the description length of the weights," in *Proceedings of the 6th Annual ACM Conference on Computational Learning Theory (COLT '93)*, pp. 5–13, July 1993.
- [45] D. J. C. MacKay, "Free energy minimisation algorithm for decoding and cryptanalysis," *Electronics Letters*, vol. 31, no. 6, pp. 446–447, 1995.
- [46] M. J. Beal, *Variational algorithms for approximate Bayesian inference*, Ph.D. thesis, University College London, London, UK, 2003.
- [47] S. Kullback and R. A. Leibler, "On information and sufficiency," *Annals of Mathematical Statistics*, vol. 22, pp. 79–86, 1951.
- [48] H. Helmholtz, "Concerning the perceptions in general," in *Treatise Physiological Optics*, J. P. C. Southall, Ed., vol. 3, Dover, New York, NY, USA, 3rd edition, 1866/1962.
- [49] R. L. Gregory, "Perceptual illusions and brain models," *Proceedings of the Royal Society of London. Series B*, vol. 171, no. 24, pp. 279–296, 1968.
- [50] R. L. Gregory, "Perceptions as hypotheses," *Philosophical transactions of the Royal Society of London Series B*, vol. 290, no. 1038, pp. 181–197, 1980.
- [51] D. Kersten, P. Mamassian, and A. Yuille, "Object perception as Bayesian inference," *Annual Review of Psychology*, vol. 55, pp. 271–304, 2004.
- [52] D. C. Knill and A. Pouget, "The Bayesian brain: the role of uncertainty in neural coding and computation," *Trends in Neurosciences*, vol. 27, no. 12, pp. 712–719, 2004.
- [53] K. Friston, K. Stephan, B. Li, and J. Daunizeau, "Generalised filtering," *Mathematical Problems in Engineering*, vol. 2010, Article ID 621670, 2010.
- [54] D. Mumford, "On the computational architecture of the neocortex—II The role of cortico-cortical loops," *Biological Cybernetics*, vol. 66, no. 3, pp. 241–251, 1992.
- [55] R. E. Kass and D. Steffey, "Approximate Bayesian inference in conditionally independent hierarchical models (parametric empirical Bayes models)," *Journal of the American Statistical Association*, vol. 407, pp. 717–726, 1989.
- [56] P. Ao, "Potential in stochastic differential equations: novel construction," *Journal of Physics A*, vol. 37, no. 3, pp. L25–L30, 2004.
- [57] R. Yuan, Y. Ma, B. Yuan, and P. Ao, "Constructive proof of global lyapunov function as potential function," <http://arxiv.org/abs/1012.2721>.
- [58] P. Ao, "Emerging of stochastic dynamical equalities and steady state thermodynamics from Darwinian dynamics," *Communications in Theoretical Physics*, vol. 49, no. 5, pp. 1073–1090, 2008.
- [59] P. A. Haile, A. Hortačsu, and G. Kosenok, "On the empirical content of quantal response equilibrium," *American Economic Review*, vol. 98, no. 1, pp. 180–200, 2008.
- [60] K. Train, *Discrete Choice Methods with Simulation*, Cambridge University Press, Cambridge, UK, 2003.

- [61] J. Maynard Smith, "Byte-sized evolution," *Nature*, vol. 355, no. 6363, pp. 772–773, 1992.
- [62] J. Maynard Smith, "The units of selection," *Novartis Foundation Symposium*, vol. 213, pp. 203–217, 1998.
- [63] E. Mayr, "Change of genetic environment and evolution," in *Evolution as a Process*, A. C. Hardy and E. B. Ford, Eds., pp. 157–180, Allen and Unwin, London, UK, 1954.
- [64] N. Eldredge and S. J. Gould, "Punctuated equilibria: an alternative to phyletic gradualism," in *Models in Paleobiology*, T. J. M. Schopf, Ed., pp. 82–115, Freeman Cooper, San Francisco, Calif, USA, 1972.
- [65] G. Matheron, *Random Sets and Integral Geometry*, John Wiley & Sons, New York, NY, USA, 1975.
- [66] H. Robbins and S. Monro, "A stochastic approximation method," *Annals of Mathematical Statistics*, vol. 22, pp. 400–407, 1951.
- [67] A. Benveniste, M. Metivier, and P. Prourier, *Adaptive Algorithms and Stochastic Approximations*, Springer, Berlin, Germany, 1990.
- [68] R. A. Rescorla and A. R. Wagner, "A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement," in *Classical Conditioning II: Current Research and Theory*, A. H. Black and W. F. Prokasy, Eds., pp. 64–99, Appleton Century Crofts, New York, NY, USA, 1972.
- [69] C. J. C. H. Watkins and P. Dayan, "Q-learning," *Machine Learning*, vol. 8, no. 3–4, pp. 279–292, 1992.
- [70] P. R. Montague, P. Dayan, C. Person, and T. J. Sejnowski, "Bee foraging in uncertain environments using predictive hebbian learning," *Nature*, vol. 377, no. 6551, pp. 725–728, 1995.
- [71] W. Schultz, P. Dayan, and P. R. Montague, "A neural substrate of prediction and reward," *Science*, vol. 275, no. 5306, pp. 1593–1599, 1997.
- [72] K. J. Friston, G. Tononi, G. N. Reeke, O. Sporns, and G. M. Edelman, "Value-dependent selection in the brain: simulation in a synthetic neural model," *Neuroscience*, vol. 59, no. 2, pp. 229–243, 1994.
- [73] J. R. Hollerman and W. Schultz, "Activity of dopamine neurons during learning in a familiar task context," *Society for Neuroscience*, vol. 22, p. 1388, 1996.
- [74] J. D. Cohen, S. M. McClure, and A. J. Yu, "Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration," *Philosophical Transactions of the Royal Society B*, vol. 362, no. 1481, pp. 933–942, 2007.
- [75] J. M. Liebman, "Anxiety, anxiolytics and brain stimulation reinforcement," *Neuroscience and Biobehavioral Reviews*, vol. 9, no. 1, pp. 75–86, 1985.
- [76] T. L. Davidson, "The nature and function of interoceptive signals to feed: toward integration of physiological and learning perspectives," *Psychological Review*, vol. 100, no. 4, pp. 640–657, 1993.
- [77] A. Alcaro, R. Huber, and J. Panksepp, "Behavioral functions of the mesolimbic dopaminergic system: an affective neuroethological perspective," *Brain Research Reviews*, vol. 56, no. 2, pp. 283–321, 2007.
- [78] S. Ishii, W. Yoshida, and J. Yoshimoto, "Control of exploitation-exploration meta-parameter in reinforcement learning," *Neural Networks*, vol. 15, no. 4–6, pp. 665–687, 2002.
- [79] W. J. Freeman, "Characterization of state transitions in spatially distributed, chaotic, nonlinear, dynamical systems in cerebral cortex," *Integrative Physiological and Behavioral Science*, vol. 29, no. 3, pp. 294–306, 1994.
- [80] I. Tsuda, "Toward an interpretation of dynamic neural activity in terms of chaotic dynamical systems," *Behavioral and Brain Sciences*, vol. 24, no. 5, pp. 793–810, 2001.
- [81] V. K. Jirsa, R. Friedrich, H. Haken, and J. A. S. Kelso, "A theoretical model of phase transitions in the human brain," *Biological Cybernetics*, vol. 71, no. 1, pp. 27–35, 1994.
- [82] V. Pasquale, P. Massobrio, L. L. Bologna, M. Chiappalone, and S. Martinoia, "Self-organization and neuronal avalanches in networks of dissociated cortical neurons," *Neuroscience*, vol. 153, no. 4, pp. 1354–1369, 2008.
- [83] M. Rabinovich, R. Huerta, and G. Laurent, "Neuroscience: transient dynamics for neural processing," *Science*, vol. 321, no. 5885, pp. 48–50, 2008.
- [84] C. van Leeuwen, "Perceptual-learning systems as conservative structures: is economy an attractor?" *Psychological Research*, vol. 52, no. 2–3, pp. 145–152, 1990.
- [85] P. Ao, D. Galas, L. Hood, and X. Zhu, "Cancer as robust intrinsic state of endogenous molecular-cellular network shaped by evolution," *Medical Hypotheses*, vol. 70, no. 3, pp. 678–684, 2008.
- [86] W. Tschacher and H. Haken, "Intentionality in non-equilibrium systems? The functional aspects of self-organized pattern formation," *New Ideas in Psychology*, vol. 25, no. 1, pp. 1–15, 2007.
- [87] C. Gros, "Cognitive computation with autonomously active neural networks: an emerging field," *Cognitive Computation*, vol. 1, no. 1, pp. 77–90, 2009.
- [88] M. Breakspear and C. J. Stam, "Dynamics of a neural system with a multiscale architecture," *Philosophical Transactions of the Royal Society B*, vol. 360, no. 1457, pp. 1051–1074, 2005.
- [89] S. L. Bressler and E. Tognoli, "Operational principles of neurocognitive networks," *International Journal of Psychophysiology*, vol. 60, no. 2, pp. 139–148, 2006.
- [90] G. Werner, "Brain dynamics across levels of organization," *Journal of Physiology Paris*, vol. 101, no. 4–6, pp. 273–279, 2007.
- [91] I. Tyukin, C. van Leeuwen, and D. Prokhorov, "Parameter estimation of sigmoid superpositions: dynamical system approach," *Neural Computation*, vol. 15, no. 10, pp. 2419–2455, 2003.
- [92] I. Tyukin, T. Tyukina, and C. van Leeuwen, "Invariant template matching in systems with spatiotemporal coding: a matter of instability," *Neural Networks*, vol. 22, no. 4, pp. 425–449, 2009.
- [93] C. Van Leeuwen, "Chaos breeds autonomy: connectionist design between bias and baby-sitting," *Cognitive Processing*, vol. 9, no. 2, pp. 83–92, 2008.
- [94] S. Nara, "Can potentially useful dynamics to solve complex problems emerge from constrained chaos and/or chaotic itinerancy?" *Chaos*, vol. 13, no. 3, pp. 1110–1121, 2003.
- [95] J. Tani, M. Ito, and Y. Sugita, "Self-organization of distributedly represented multiple behavior schemata in a mirror system: reviews of robot experiments using RNNPB," *Neural Networks*, vol. 17, no. 8–9, pp. 1273–1289, 2004.
- [96] J. M. Herrmann, K. Pawelzik, and T. Geisel, "Self-localization of autonomous robots by hidden representations," *Autonomous Robots*, vol. 7, no. 1, pp. 31–40, 1999.
- [97] A. Wald, "A new formula for the index of cost of living," *Econometrica*, vol. 7, pp. 319–331, 1939.
- [98] D. S. Jones, *Elementary Information Theory*, Clarendon Press, Oxford, UK, 1979.

- [99] T. Basar and G. J. Olsder, *Dynamic Noncooperative Game Theory*, Academic Press, London, UK, 2nd edition, 1995.
- [100] E. Todorov, "Linearly-solvable Markov decision problems," in *Advances in Neural Information Processing Systems*, vol. 19, pp. 1369–1376, MIT Press, Boston, Mass, USA, 2006.
- [101] G. A. Rummery and M. Niranjan, "On-line Q-learning using connectionist systems," Tech. Rep. CUED/F-INFENG/TR 166, Engineering Department, Cambridge University, Cambridge, UK, 1994.
- [102] B. Sallans and G. E. Hinton, "Reinforcement learning with factored states and actions," *Journal of Machine Learning Research*, vol. 5, pp. 1063–1088, 2004.