

Article

Modules or Mean-Fields?

Thomas Parr * , Noor Sajid and Karl J. Friston

Wellcome Centre for Human Neuroimaging (UCL), London WC1N 3AR, UK; noor.sajid.18@ucl.ac.uk (N.S.); k.friston@ucl.ac.uk (K.J.F.)

* Correspondence: thomas.parr.12@ucl.ac.uk

Received: 1 April 2020; Accepted: 12 May 2020; Published: 14 May 2020



Abstract: The segregation of neural processing into distinct streams has been interpreted by some as evidence in favour of a modular view of brain function. This implies a set of specialised ‘modules’, each of which performs a specific kind of computation in isolation of other brain systems, before sharing the result of this operation with other modules. In light of a modern understanding of stochastic non-equilibrium systems, like the brain, a simpler and more parsimonious explanation presents itself. Formulating the evolution of a non-equilibrium steady state system in terms of its density dynamics reveals that such systems appear on average to perform a gradient ascent on their steady state density. If this steady state implies a sufficiently sparse conditional independency structure, this endorses a mean-field dynamical formulation. This decomposes the density over all states in a system into the product of marginal probabilities for those states. This factorisation lends the system a modular appearance, in the sense that we can interpret the dynamics of each factor independently. However, the argument here is that it is *factorisation*, as opposed to *modularisation*, that gives rise to the functional anatomy of the brain or, indeed, any sentient system. In the following, we briefly overview mean-field theory and its applications to stochastic dynamical systems. We then unpack the consequences of this factorisation through simple numerical simulations and highlight the implications for neuronal message passing and the computational architecture of sentience.

Keywords: stochastic dynamics; modularity; density dynamics; message passing; Bayesian mechanics

1. Introduction

Attempts to understand neuroanatomical and psychological organisation have often appealed to the notion of a ‘module’ [1–5]. The basic idea is that cognition depends upon a set of specialised modules that operate (almost) independently of one another. Each module is thought to receive a specialised form of input—often a specific sensory modality—and provides a low dimensional output to other modules. It is easy to see the appeal of this kind of formulation. Just as we think of the heart as an organ to pump blood, the kidneys to filter it, and the lungs to oxygenate it, the modular perspective on cognitive function lets us (literally) organize the brain into constituent organs that each play their own role in processing information. The occipital cortices are ‘for’ processing visual data, the ventral visual stream ‘for’ identifying the thing that caused these data and the dorsal stream ‘for’ locating these causes. Often, this teleological perspective is motivated in terms of evolutionary psychology [6]. Pragmatically, this suggests an approach to evaluating cognitive function. If we can think of the brain in terms of functionally specialised modules, it should be possible to design experiments and cognitive tests that interrogate these, independently of one another. In this paper, we argue that the emergence of a modular architecture is more simply expressed in terms of factorisation. This perspective arises from an approach developed in statistical physics called mean-field theory [7,8]. The basic idea is that a probability distribution over the components of a system may be approximated by the product of the distributions for each component (or groups of components) of that system. This treatment assumes

we can treat parts of the system as operating independently to other parts, just as modules are treated as independent of one another. In addition to neurobiology [9,10], applications of mean-field theory are broad, and have been used to find tractable solutions to problems in fields as diverse as statistics [11], soft-matter physics [12], epidemiology [13], game theory [14], and financial modelling [15].

Section 2 provides a review of mean-field theory, the problem it was developed to solve, and the form of the solution. Interestingly, this solution does not involve complete independence of each factor. Instead it ensures the components of a system depend upon one another via their mean-fields—so-called because only the average values of other components matter. Section 3 takes the concept of mean-field theory and places it in a dynamical context. We set out the density dynamics of mean-field systems at (non-equilibrium) steady state. Doing so reveals that each factorised component appears to undergird its own steady state density. However, the steady state density of each factor depends upon the mean field of other factors. Section 4 introduces some minimal simulations that aim to build an intuition as to how this works in practice. These numerical analyses are designed to illustrate the ideas introduced in earlier sections as simply as possible. Here, we see a simple form of functional (modular) specialisation, and the emergence of a separation of timescales that is characteristic of hierarchical neuronal dynamics. Section 5 highlights the link between mean-field formalism, inference, and the message passing between populations of neurons. This rests on the fact that the simplest tractable way to make inferences about the causes of (sensory) data is to use a mean field approximation that underwrites a form of variational or approximate Bayesian inference. In short, we can study the properties of stochastic dynamical systems—like the brain—through mean-field assumptions. This should not be interpreted as a model of the brain—rather it is an approach to understanding stochastic systems with sparse dependency structures, of which the brain is a paradigmatic example. We suggest that accounts of cognitive function in terms of modular architectures rest upon an intuitive application of mean-field theory. Making this explicit provides a useful perspective on brain function and lets us exploit established tools from stochastic physics. We start with an overview of these tools.

2. Mean-Field Theory

The origins of mean-field theory are in physics [7,8]. They were invoked to study systems described by a Gibbs' measure. This is an expression of the statistical properties of a system that says that the probability density of a system being in a particular state x decreases as the energy associated with that state increases. In other words, the higher the total energy of the system in each configuration, the lower the probability of that configuration. Turing this on its head, energy may be thought of as a measure of the improbability of a configuration. For reasons that will be clearer later, we are interested in systems with a second random variable that takes the value y . This is a parameter can change the shape of the energy landscape for x . In the context of the neurosciences, x could indicate (log) neuronal firing rates with y indicating sensory stimulation. Through Bayes' theorem, we can interpret the variables in this system in terms of joint, conditional, and marginal probability densities:

$$\begin{aligned}
 p(x|y) &= \frac{1}{Z(y)} e^{-\beta \mathcal{H}(x,y)} \\
 Z(y) &\triangleq \int_{-\infty}^{\infty} e^{-\beta \mathcal{H}(x,y)} dx \\
 &\Rightarrow \ln p(x|y) = -\beta \mathcal{H}(x,y) - \ln Z(y) \\
 &\Rightarrow \begin{cases} \ln p(x,y) = -\beta \mathcal{H}(x,y) \\ \ln p(y) = \ln Z(y) \end{cases}
 \end{aligned} \tag{1}$$

The total energy of the system is given by the Hamiltonian (\mathcal{H}). For classical dynamical systems, this is a scalar function. For quantum dynamical systems, this is a linear operator whose eigenvalues are interpretable as energies. Equation (1) is more general than it appears at a first glance. While the expression in the first line may seem restrictive, the Gibbs' form in the first equality of Equation (1) can be used to express any exponential family probability distribution by choosing different forms for the

Hamiltonian. Some common examples are given in Table 1. The integral in the second equality must be replaced by a sum when the support of the distribution is categorical. The β parameter is sometimes referred to as an ‘inverse temperature’ parameter, as it is inversely proportional to the temperature of a physical system. This determines how ‘peaky’ the distribution is, with high β concentrating probability mass on a small region of space, and low β leading to a more even distribution of probability mass.

Table 1. Exponential family distributions.

Distribution	Support	Hamiltonian
Gaussian	$x \in \mathbb{R}$	$\frac{1}{2\beta}(x - \mu) \cdot \Pi(x - \mu)$
Multinomial ¹	$x_i \in \{0 \dots N\}$ $i \in \{1, \dots, K\}$ $\sum_i x_i = N$	$-\frac{1}{\beta} \sum_i x_i \ln d_i$
Dirichlet ²	$x_i \in (0, 1)$ $i \in \{1, \dots, K\}$ $\sum_i x_i = 1$	$\frac{1}{\beta} \sum_i (1 - \alpha_i) \ln x_i$
Gamma	$x \in (0, \infty)$	$\frac{1}{\beta}(bx + (1 - a) \ln x)$

¹ Special cases include Categorical ($K > 2, N = 1$), Binomial ($K = 2, N > 1$), and Bernoulli ($K = 2, N = 1$) distributions.

² A special case is the Beta distribution ($K = 2$).

The denominator—or normalising constant—(Z) of the first line of Equation (1) is an important quantity in thermodynamics called a partition function. This is closely related to another quantity called Helmholtz free energy [17,18]:

$$\begin{aligned}
 F(y) &\triangleq -\frac{1}{\beta} \ln Z(y) \\
 &= \frac{1}{\beta} \ln p(x|y) + \mathcal{H}(x, y) \\
 &= E_{p(x|y)} \left[\frac{1}{\beta} \ln p(x|y) + \mathcal{H}(x, y) \right] \\
 &= U(x, y) - TS(x, y) \\
 U(x, y) &\triangleq E_{p(x|y)} [\mathcal{H}(x, y)] \\
 S(x, y) &\triangleq -\frac{1}{\beta T} E_{p(x|y)} [\ln p(x|y)]
 \end{aligned} \tag{2}$$

Here, E indicates an expectation (i.e., average), U is the internal energy of the system, T is its temperature, and S is its entropy. The third equality rests upon the fact that the Helmholtz free energy does not depend upon x , so:

$$\begin{aligned}
 E_{p(x|y)} [F(y)] &= F(y) \\
 \Rightarrow E_{p(x|y)} \left[\frac{1}{\beta} \ln p(x|y) + \mathcal{H}(x, y) \right] &= \frac{1}{\beta} \ln p(x|y) + \mathcal{H}(x, y)
 \end{aligned} \tag{3}$$

With these preliminaries in place, we are now able to define the problem for which mean-field theory is the solution. This problem arises when we know only the Hamiltonian. Simply put, the partition function (Z) is hard to compute. This is due to the difficulty of calculating the integral in Equation (1) for all but the simplest Hamiltonians. Without the partition function, we cannot calculate the conditional density of x given y . The mean-field approach starts by considering a simpler (reference) system, where there are no interactions between the constituents of the system. This absence of interactions is known as a mean-field assumption:

$$\begin{aligned}
 q(x|y) &= \prod_i q(x_i|y) \\
 q(x_i|y) &= \frac{e^{-\beta h_i(x_i,y)}}{Z_i(y)} \\
 \mathcal{H}_q(x, y) &= \sum_i h_i(x_i, y)
 \end{aligned} \tag{4}$$

This system factorises into a series of marginal distributions in virtue of the decomposition of the Hamiltonian into a sum of Hamiltonians for each component of the system. We refer to the distribution q as a variational density [19]. At this point, we can appeal to the Bogolyubov inequality [20]. This is a special case of Jensen's inequality that says that the Helmholtz free energy of the interacting system is always less than if we calculated the free energy using the original Hamiltonian but replace the conditional probability with the variational density. We refer to the latter as the variational free energy and use the subscript q to distinguish this from the Helmholtz free energy. Re-expressing in terms of a Kullback-Leibler (KL) Divergence (A relative entropy (the average log ratio of two densities) that is always greater than or equal to zero (by Jensen's inequality)) the Bogolyubov inequality becomes clear:

$$\begin{aligned}
 F_q(y) &\triangleq \mathbb{E}_{q(x|y)} \left[\frac{1}{\beta} \ln q(x|y) + \mathcal{H}(x, y) \right] \\
 &= \mathbb{E}_{q(x|y)} \left[\frac{1}{\beta} \ln q(x|y) - \frac{1}{\beta} \ln p(x|y) \right] - \frac{1}{\beta} \ln Z(y) \\
 &= \frac{1}{\beta} \underbrace{D_{KL}[q(x|y)||p(x|y)]}_{\geq 0} + F(y) \\
 &\geq F(y)
 \end{aligned} \tag{5}$$

Equation (5) says that the variational free energy (F_q) is an upper bound on the Helmholtz free energy (F). The implication is that, by minimising the latter, we should arrive at a good approximation of the former. This converts the difficult integration problem of Equation (1) into a much easier optimisation problem. Variational approaches of this sort have a long history, perhaps most famously in the formulation of quantum mechanics in terms of distributions over alternative paths a particle might follow [21]. Crucially, the factorisation of the variational density means we can optimise each factor independently. It is this property that lends a modular aspect to particular kinds of random dynamical system.

To understand how the different factors interact, it is worth highlighting that the Hamiltonian of the interacting system can itself be decomposed into a sum of factors. These will not be the independent factors of the mean-field reference system. Instead, they are conditional probability densities. Many elements in the sum are functions of more than one component of the system, and each component can contribute to more than one factor. Figure 1 illustrates a graphical notation used to represent the decomposition of a Hamiltonian. This general formalism has been exploited in signal processing [22], Newtonian [23] and quantum dynamics [24], and neurobiology [25,26]. Each square factor indicates a potential (φ_K) whose argument (x_K) is some subset of x from a region (K) of the graph involved in that potential. For example, region 6 includes (x_1, x_6) , as these are the variables linked to the φ_6 node. Crucially, regions overlap such that x_6 participates in regions 6 and 7. The Hamiltonian is given by the sum of these potentials:

$$\mathcal{H}(x, y) = \sum_K \varphi_K(x_K, y) \tag{6}$$

In general, not every potential will include y as an argument. In the example of Figure 1, only factors 18, 19, 20, and 21 include y as an argument, and each of these includes a different subset of the y variables. The Hamiltonian is constructed with three things in mind. The first is simplicity. To ensure this, we have used quadratic potentials that simplify the treatment of density dynamics in Section 3. The second is sparsity, which is a characteristic feature of brain-like systems. Sparsity means that each component of a system (e.g., neuron in a brain) interacts directly with relatively few other components. The third is that there are several different points at which the y variables may influence the system. This is consistent with alternative sensory modalities in nervous systems. We can now express the solution to the problem of finding the partition function as follows:

$$\begin{aligned}
 q(x_i|y) &= \arg \min_{q(x_i|y)} F_q(y), \forall i \\
 &\Leftrightarrow h_i(x_i, y) = \sum_{\{K: x_i \in x_K\}} E_{q_{K \setminus i}}[\varphi_K(x_K, y)] \\
 &\Rightarrow F_q(y) \approx F(y) \\
 &\Rightarrow q(x|y) \approx p(x|y) \\
 q_{K \setminus i} &\triangleq \frac{q(x_K|y)}{q(x_i|y)}
 \end{aligned} \tag{7}$$

The approximate equality between the ‘ p ’ and ‘ q ’ distributions rests upon the assumption that the latter comprises a product of marginal factors (Equation (4))—which is not assumed for the former. The quality of the approximation may be quantified by the (negative) KL-Divergence between the two. Note that this is exactly the bound that appears in Equation (5) quantifying the difference between the associated partition functions. This accounts for the implication in Equation (7) that, when the partition functions are approximately equal, the KL-Divergence is approximately zero, and the ‘ p ’ and ‘ q ’ distributions are approximately equal. The second line expresses the ‘mean-field’—the average of the local potentials. There are two important things to draw from Equation (7). First, for the mean field associated with a factor, only the average values of the other factors matter. Second, we can ignore most of the terms in the sum of potentials comprising the original Hamiltonian. We only need those potentials in which our variable of interest participates, i.e., the ‘local’ potentials. This will become important in Section 5, where we revisit this idea in relation to the sparse connectivity structure of the brain [25].

While outside the scope of this paper, there are generalisations of mean-field theory that rely upon more sophisticated choices for the variational distribution. Cluster variational methods [27,28], based upon Kikuchi free energies, offer a much more general formulation. In brief, these employ a reference system with overlapping factors, corrected for the overlaps. It is the presence of these overlaps that distinguishes such approaches from mean-field theory, which is predicated upon the absence of overlaps. Table 2 sets out the form of the Hamiltonian associated with the variational distributions for a few key examples. Each of these is associated with a different inference scheme that minimises the associated variational free energy.




Figure 1. This schematic illustrates how the decomposition of a Hamiltonian into the sum of potentials may be represented graphically. This is a factor graph that represents each potential as a square node. The arguments of each potential are represented as circles connected to that square node. The y arguments of the potentials are represented as smaller squares. The arrows on some of the edges inherit from the interpretation of potentials as log conditional probabilities. If a random variable A is conditionally dependent on a variable B , the factor linking the two will include an arrow pointing towards A . The factor graph shown here is the (arbitrarily constructed) Hamiltonian that we will employ in the simulations in subsequent figures. This assumes a quadratic form for each potential. The details of these potentials are not important and could be replaced with any alternative quadratic functions. For readers interested in the precise formulation used in the simulations that follow, please see the Matlab routines referred to in the software note. In brief, each potential is centred upon a linear function of the mode of the neighbouring potential. An important feature of this structure is the sparsity of conditional dependencies. Each factor connects at most two variables. We assume $x_i \in \mathbb{R}^2$ in what follows. Uppercase subscripts are used to identify larger groups of x (i.e., $x_K \in \mathbb{R}^{\geq 2}$), corresponding to the argument of a given potential.

Table 2. Variational distributions.

Name	Hamiltonian	Comments
Mean-field	$\sum_i h_i(x_i, y)$	As in the main text, x is divided into non-overlapping subsets (x_i), each of which is associated with its own Hamiltonian. The inference scheme associated with this approximation is known as <i>Variational message passing</i> [11].

Table 2. Cont.

Name	Hamiltonian	Comments
Bethe	$\sum_{ij} h_{ij}^{(2)}(x_i, x_j, y) - \sum_k (c_k^{(1)} - 1) h_k^{(1)}(x_k, y)$	This expression uses a series of overlapping pairwise (superscript 2) Hamiltonians, that are then ‘corrected’ for these overlaps by subtracting singleton (superscript 1) Hamiltonians. Here, c_k is the number of pairwise factors that include x_k as an argument. The inference scheme associated with this approximation is known as (<i>loopy</i>) <i>Belief propagation</i> [29].
Kikuchi	$\sum_R c_R^{(i)} h_R^{(i)}(x_R^{(i)}, y)$ $c_R^{(i)} \triangleq 1 - \sum_{\{K:R \subset K\}} c_K^{(i+1)}$	This expression generalises the above approximations. Here, the subscripts index regions, while the superscript indexes the size of that region. In this expression, $x_R^{(i)}$ includes all elements of x in region R at scale i . Here, regions may overlap. If all regions are of size 1, this reduces to a mean-field approximation. If some are size 1 and others size 2, this is the Bethe approximation. Inference schemes based on the Kikuchi approximation are known as <i>Cluster variational methods</i> or <i>Generalised belief propagation</i> [27,28].

3. Non-Equilibrium Stochastic Dynamics

In this section, we take a step back and think about the dynamics of stochastic systems subject to the analyses of the previous section. These are systems that have attained a (possibly non-equilibrium) steady state, in the sense that the Hamiltonian is interpretable as a (static) log probability density. The first step in understanding what this means is to note that there are multiple equivalent ways in which the dynamics of a stochastic system may be formulated. We will focus upon two of these. One is a stochastic differential equation, which expresses equations of motion that depend upon a deterministic flow (f) and random fluctuations (ω). We will assume these fluctuations are normally distributed and uncorrelated over time or space. The second formulation we appeal to is afforded by a Fokker–Planck equation (a.k.a., a Kolmogorov forward equation). Instead of dealing with specific instances of a random system, Fokker–Planck equations deal with the dynamics of their probability density [30]. For concision throughout, we will use the dot notation to indicate partial time derivative:

$$\left. \begin{aligned} \dot{x} &= f(x, y) + \omega \\ E[\omega(\tau) \cdot \omega(t)] &= 2\Gamma\delta(\tau - t) \end{aligned} \right\} \Leftrightarrow \dot{p}(x|y) = \nabla_x \cdot ((\Gamma \nabla_x - f(x, y))p(x|y)) \tag{8}$$

In Equation (8), the amplitude of the random fluctuations is given by a diffusion tensor (2Γ). The δ -symbol indicates a Dirac delta function that ensures the covariance of the fluctuations at two time points (t and τ) is zero unless these times coincide, i.e., the fluctuations are temporally uncorrelated (c.f., a Wiener process). The Fokker–Planck equation on the right shows the rate at which probability mass enters or leaves an infinitesimally small region of space around x . Appendix A introduces the Fokker–Planck equation and links it to the stochastic differential equation on the left. However, the intuition is relatively simple. Imagine a drop of ink in water. Initially, the distribution of ink has a very sharp peak as it is concentrated in one place. This implies a large negative second derivative at this point, and relatively fast dispersion of the ink. As this peak is dispersed, and the second derivative becomes closer to zero, the rate at which ink leaves the initial location reduces. If the amplitude of fluctuations is greater (e.g., the water is boiling), the ink will spread out faster. This accounts for the term weighted by the diffusion tensor. The intuition for the role of the deterministic flow (f) is simpler. If there are currents in the water, the ink will leave those regions with fast flowing currents faster than regions of slower currents. The gradient of the current is key, as a positive gradient implies the currents leaving a region are faster than those entering it, while negative implies the opposite.

Using Equation (8), and the assumption that the rate of change of the probability density is zero when described by the Gibbs' measure of Section 2, we can find an expression for the equations of motion in terms of the gradients of the Hamiltonian [31,32]:

$$\begin{aligned}
 p(x|y) \propto e^{-\beta\mathcal{H}(x,y)} &\Leftrightarrow \dot{p}(x|y) = 0 \\
 \Rightarrow \nabla_x \cdot \left((\Gamma \nabla_x - f(x,y)) e^{-\beta\mathcal{H}(x,y)} \right) &= 0 \\
 \Rightarrow \Gamma \nabla_x e^{-\beta\mathcal{H}(x,y)} - f(x,y) e^{-\beta\mathcal{H}(x,y)} &= Q \nabla_x e^{-\beta\mathcal{H}(x,y)} \\
 \Rightarrow f(x,y) = -\beta(\Gamma - Q) \nabla_x \mathcal{H}(x,y) &
 \end{aligned} \tag{9}$$

The matrix Q is defined such that all its eigenvalues are pure imaginary or zero, ensuring the term on the right-hand side of the third line is divergence free. For the purposes of this paper, we will assume a block diagonal form for Q , where each matrix on the diagonal is a square, skew-symmetric matrix of dimension 2. We have assumed in the above that neither Q nor Γ vary with x . However, a more general form for Equation (9) can be constructed that allows these to vary [33]. The first and second rows of plots in Figure 2 show what happens when we simulate this system, by substituting the final line of Equation (9) into the stochastic differential equation in Equation (8). The temperature parameter (β) for these simulations is set at one. The dispersion of the steady-state density is therefore determined solely by the Hessian of the Hamiltonian. While simulating a single instantiation of these dynamics leads to a very noisy trajectory (first row of plots), simulating multiple instances and averaging reveals the self-organisation of this system into an 'x' shape.

While we could keep adding additional instances to this simulation and get incremental improvements to the characterisations of the distribution, a simple approach is to simulate the density dynamics directly. This gives us the results we would have found in the limit of infinitely many simulations of specific instances. The difficulty with this is that Fokker–Planck equations using the flow from Equation (9) directly involve an unwieldy covariance matrix for systems comprising many particles. However, we can simplify this problem by solving for individual factors. Applying the mean-field approximation from the previous section, we find the simpler expression:

$$\begin{aligned}
 q(x_i|y) \propto e^{-\beta h_i(x_i,y)} &\Leftrightarrow \dot{q}(x_i|y) = 0 \\
 \Rightarrow f_i(x,y) = -\beta(\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i,y) &
 \end{aligned} \tag{10}$$

Note that this is only a function of x_i and y , and no other system components. There is a sense in which this can be interpreted as 'information encapsulation' [34], one of the key features ascribed to modular architectures. Substituting this into the Fokker–Planck equation, we have:

$$\dot{q}(x_i|y) = \nabla_{x_i} \cdot \left((\Gamma_{ii} \nabla_{x_i} + \beta(\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i,y)) q(x_i|y) \right) \tag{11}$$

There are several methods available for solving Equation (11). Broadly, these include discretising over space or assuming a functional form for the probability density. For the former, this means integrating for each pixel (in two dimensions) based upon numerical gradients and Laplacians. The latter involves solving for the associated parameters. Either approach may be used here. We adopt the latter, which has the advantage of requiring fewer dimensions than a discretisation-based approach. Re-expressing this in terms of the sufficient statistics of the probability density—its mean and variance—we have (see Appendix B):

$$\begin{aligned}
\dot{\mu}_i &= -\beta(\Gamma_{ii} - Q_{ii})E_{q(\mu_i|y)}[\nabla_{x_i}h_i(x_i, y)] \\
&\approx -\beta(\Gamma_{ii} - Q_{ii})\sum_{\{K:x_i \in x_K\}}\left(\nabla_{x_i}\varphi_K(x_K, y)\Big|_{x_K=0} + \nabla_{x_i x_K}\varphi_K(x_K, y)\Big|_{x_K=0}\mu_K\right) \\
\dot{\Sigma}_{ii} &= 2\Gamma_{ii} \\
&\quad -\beta E_{q(x_i|y)}[\Delta x_i \nabla_{x_i}h_i(x_i, y)^T](\Gamma_{ii} - Q_{ii})^T \\
&\quad -\beta(\Gamma_{ii} - Q_{ii})E_{q(x_i|y)}[\nabla_{x_i}h_i(x_i, y)\Delta x_i^T] \\
&\approx 2\Gamma_{ii} \\
&\quad -\beta\sum_{\{K:x_i \in x_K\}}\nabla_{x_i x_i}\varphi_K(x_K, y)\Big|_{x_K=0}\Sigma_{ii}(\Gamma_{ii} - Q_{ii})^T \\
&\quad -\beta(\Gamma_{ii} - Q_{ii})\sum_{\{K:x_i \in x_K\}}\nabla_{x_i x_i}\varphi_K(x_K, y)\Big|_{x_K=0}\Sigma_{ii}
\end{aligned} \tag{12}$$

The first of these equations sets out the dynamics of Equation (10)—the expected rate of change—under quadratic assumptions about the form of the Hamiltonian. For the simulations reported here, these assumptions hold by construction of the Hamiltonian as a quadratic function. More generally, this assumption depends upon local Taylor series approximations of the Hamiltonian. The second equation gives the dynamics of the covariance. Note that this equation is zero when the covariance is equal to the inverse of the sum of Hessian matrices (up to a scale factor β). The dynamics of the covariance provide an interesting perspective on the change in entropy of the system over time. Specifically, Equation (12) indicates that the rate of change of the covariance may be positive or negative. Remembering that the entropy of a normal distribution depends only on the covariance (and not the mode), we see that the system may increase or decrease its entropy. Consistent with the fluctuation theorems of stochastic thermodynamics [35], this highlights that the direction of change in entropy depends upon whether the initial or steady-state density has the greater dispersion. The third row of Figure 2 shows the results when Equation (12) is used to simulate the density dynamics of a system with the Hamiltonian of Figure 1. In Section 4, we unpack these dynamics in relation to modular theories.

Equation (12) can be simplified by introducing auxiliary variables (Π, ε):

$$\begin{aligned}
\dot{\mu}_i &\approx -\beta(\Gamma_{ii} - Q_{ii})\sum_{\{K:x_i \in x_K\}}\Pi_i^K \varepsilon_i^K \\
\dot{\Sigma}_{ii} &\approx 2\Gamma_{ii} - \beta\sum_{\{K:x_i \in x_K\}}\Pi_i^K \Sigma_{ii}(\Gamma_{ii} - Q_{ii})^T - \beta(\Gamma_{ii} - Q_{ii})\sum_{\{K:x_i \in x_K\}}\Pi_i^K \Sigma_{ii}
\end{aligned} \tag{13}$$

The auxiliary variables are defined as follows:

$$\begin{aligned}
\Pi_i^K &\triangleq \nabla_{x_i x_i}\varphi_K(x_K, y)\Big|_{x_K=0} \\
\varepsilon_i^K &\triangleq \mu_i - \eta_i^K(\mu_{K \setminus i}) \\
\eta_i^K(\mu_{K \setminus i}) &\triangleq -(\Pi_i^K)^{-1}\left(\nabla_{x_i x_{K \setminus i}}\varphi_K(x_K, y)\Big|_{x_K=0}\mu_{K \setminus i} + \nabla_{x_i}\varphi_K(x_K, y)\Big|_{x_K=0}\right)
\end{aligned} \tag{14}$$

Equations (13) and (14) provide a useful intuition as to the behaviour of the system. It implies that the mode of each marginal density changes such that it minimises the difference (ε) between itself and a ‘target’ value (η), where the latter is a function of the modes of the other marginals with which it shares a potential. Each mode effectively chases (or ‘tracks’) a moving target until all modes have reached their attracting points. This mediates a form of synchronisation, on average, between the factorised components of the system. However, this does not mean the marginals contain information about the joint densities. Instead, interactions are mediated via the mean-fields as in Equation (7).

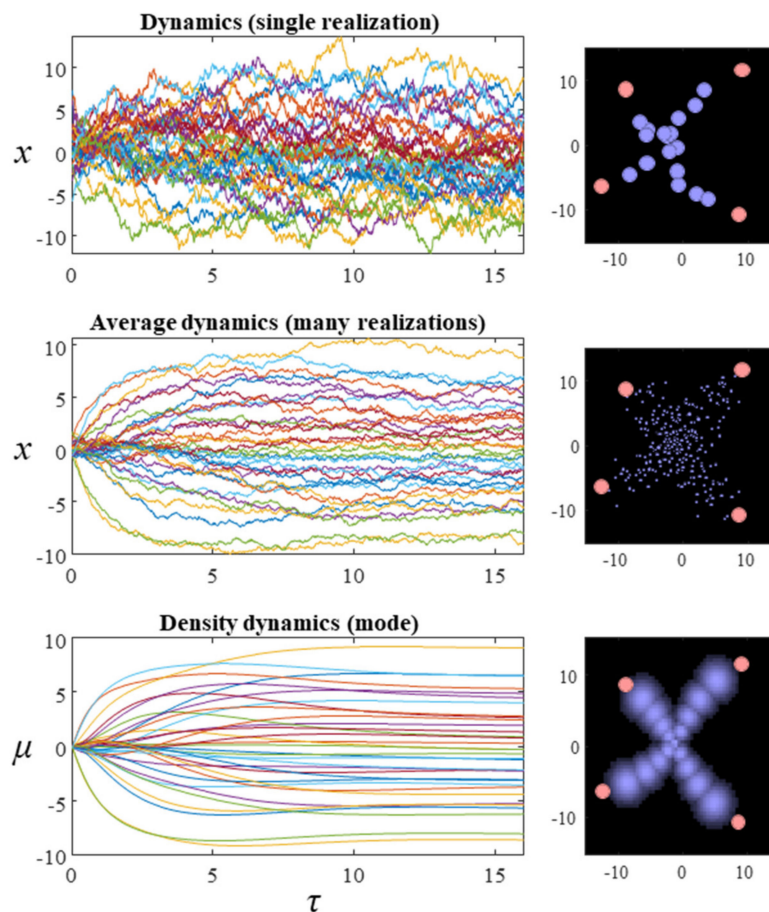


Figure 2. The plots in this figure illustrate the evolution of the random dynamical system whose Hamiltonian is shown in Figure 1. The plots on the left show the evolution of the system over time. This is a 34-dimensional system, which is shown on the right in terms of 17 particles, whose positions are described by two coordinates. The plots on the right show the final configuration at the end of the simulation. The first row shows a single realization of a stochastic trajectory. The second averages over 16 realizations of the trajectory. The third row shows the density dynamics under a Laplace approximation. The mean-field factorisation treats each particle independently (so each factor is a bivariate normal distribution). The filled pink circles in the plots on the right illustrate the values of the y variables (which are fixed). For ease of visibility, the intensity of each of the densities superimposed on this image have been normalised, such that their mode is the same intensity (regardless of the probability density at that mode).

This treatment may sound very abstract and technical, however, it forms the basis for much of physics as we know it. Furthermore, it has enormous practical implications. Effectively, the simulations in Figure 2 show that it is possible to create highly structured ensemble dynamics (here a nonlinear 17-body problem with random fluctuations) with a desired ‘shape’. In other words, we can effectively write down a probabilistic description of what we want a system to look like, and then use the mean field approximation to realise that kind of system. In engineering, this would be known as directed self-assembly and is a central part of nanotechnology [36,37]. In the neurosciences, the (dynamic causal) modelling of neural interactions rests upon the mean field approximation in Equation (14) [38], creating a distinction between neural mass and mean field models [39].

The different perspectives on the same underlying dynamics shown in Figure 2 provide an interesting point of connection to different kinds of probabilistic inference scheme used widely in statistics and machine learning. Broadly, approximate inference techniques are divided into two classes. The first relies upon sampling, and include Markov Chain Monte Carlo (MCMC) approaches such as

the Metropolis-Hastings algorithm [40] or Gibbs' sampling [41,42]. Special cases of MCMC, including the Metropolis-adjusted Langevin algorithm [43] are based upon the dynamics given by Equation (9) to ensure a target distribution is attained after sufficient time. A more general form of Equation (9) has been used explicitly in constructing MCMC samplers [33]. The second approach is to work directly with the density dynamics by assuming a parameterised form for the density and optimising these parameters [19], i.e., variational inference. The first two rows of plots in Figure 2 can be thought of as showing how sampling approaches progress, while the third row is an example of a variational scheme.

4. Factors and Modules

In this section, we draw upon the stochastic dynamics of Figure 2. The Hamiltonian that underwrites this specifies a pattern in which the location of each component of the system is conditionally dependent upon locations of other components. Our first step is to note that we can look at each of these components independently. Figure 3 shows the trajectory of the modes, and the final density, for each factorised density under the mean-field approximation. The reciprocal dependencies between these factors (i.e., the mean-fields) are shown as arrows. Note the spiral trajectories. These result from the combination of solenoidal and curl-free flows (down and around the gradients of the Hamiltonian, respectively). The decomposition of a single system into a series of interacting subsystems offers our first hint at 'modularity'.




Figure 3. This figure decomposes the density dynamics of Figure 2 in line with the mean-field partition. The arrows here indicate the influence of each marginal density on another via their associated mean-fields. In other words, they represent the non-zero elements of the Jacobian for the vector μ , with elements μ_i , whose rate of change is given in Equation (13). Each image shows the probability density at the end of the simulation (in blue) and the trajectory of the mode throughout the simulation (white). Note the highly precise distribution over the central factor, which is constrained by its four neighbours. The key message to take away from this figure is that the mean-field approximation separates the full system of Figure 2 into a series of smaller systems that influence one another only through their averages.

The next stage in our analysis is to think about the consequences of perturbing the system, to see how each marginal density responds. We can do this by interpreting y as sensory data and manipulating variables. This resembles standard approaches in neuroscience that measure the brain's response to experimental sensory stimuli. Figures 4 and 5 show what happens when we perturb the upper right (y_1 in Figure 1) sensory input, the lower left (y_2 in Figure 1) input, or both.

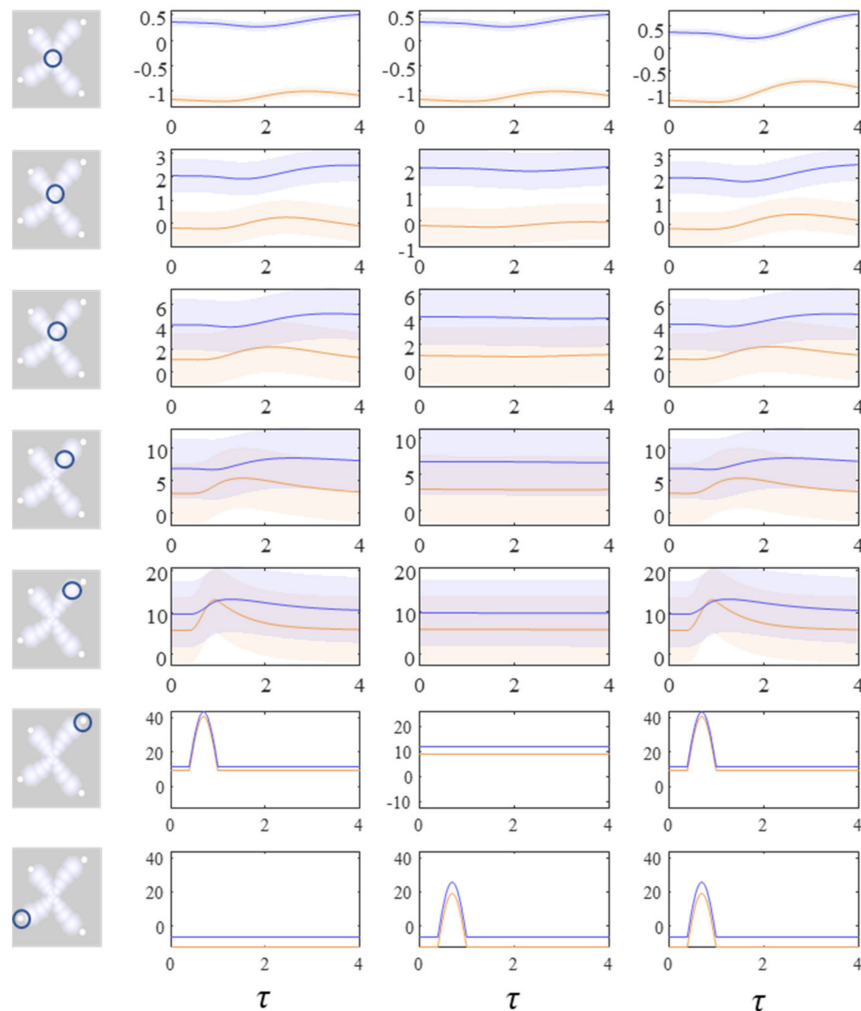


Figure 4. These plots show the consequences of perturbing the y variables on the density dynamics (depicted as the mode and surrounding 90% credible intervals) for each factor. The images on the left indicate which factor is shown in each row of the plots. Each column of plots shows a separate simulation in which different perturbations are applied to the y variables. In this figure, the plots show the central factor (first row) through to the upper right factor (fifth row). The lower two rows show the y variables in the upper right and lower left. These are perturbed by introducing a sinusoidal impulse. The first column of plots shows the response to the upper right perturbation. The second column shows the limited response to the lower left perturbation. The third column shows the increased recruitment of more central regions in the presence of both perturbations. The key point to take away from this figure is that a simple form of ‘information encapsulation’ or functional specialisation occurs in the extremities, with specific responses to, and only to, one sort of y variable. Over a hierarchy of timescales and progressively prolonged responses evocative of delay-period firing in working memory tasks, the factors become progressively multimodal. Figure 5 shows the lower left factors in the same simulation.

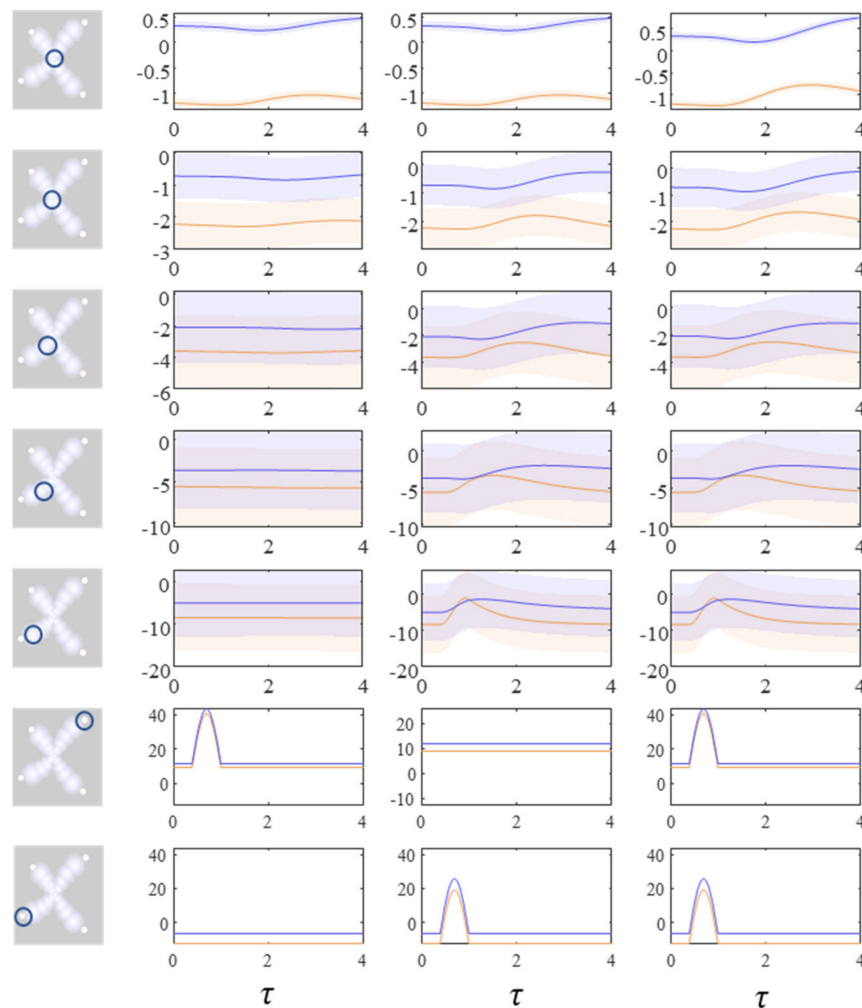


Figure 5. These plots complement those of Figure 4, illustrating the same perturbations and their consequences for the lower left modules. Here, there is little effect of the upper right y perturbation until we reach more central regions. However, there is a response to the lower left y perturbation that was not seen in Figure 4. For details on the format of these plots, please see the legend of Figure 4.

Figure 4 shows the density dynamics of the central and upper right factors (see Figure 3), while Figure 5 shows these for the central and lower left factors. There are three things to take away from these plots. First, they illustrate a form of functional specialisation, in that the lower left factors respond to changes in the lower left stimulus but not to the upper right stimulus, and vice versa for the upper right factors. In other words, we have segregated sensory streams that deal with different aspects of the sensorium: similar to cognitive processing associated with visual [44], auditory [45], language [46], and temporal [47] tasks. The second thing to note is that the timescale of the responses is slower (peaking later and persisting longer) the closer to the central factor (x_1). This mimics the (slow and fast) temporal separation seen in neurobiological hierarchies [48–51]. It also implies a simple form of working memory, in the sense that the effects of the stimulus persist long after it has been removed. Finally, the more central factors respond to both sensory inputs, and show a greater response when both are presented simultaneously. Here, we have evidence in favour of multimodal factors analogous to those brain cells that respond to stimuli presented to different sensory modalities [52–56]. Multimodal properties of this sort speak to the importance of functional integration alongside modular segregation [2,57], heightened during cognitive processing [58].

5. Neuronal Message Passing

In the preceding sections, we used an arbitrarily constructed random dynamical system to illustrate a factorised (or modularised) account of systems with a sparse conditional independency structure at steady state. The resulting density dynamics show a form of functional segregation with distinct ‘sensory’ streams. As these converge upon one another, we see the emergence of a simple form of multisensory integration, based upon the expectation values of the sensory streams. Along these streams, each factor operates with a different temporal scale, much as sequences of cortical regions in sensory hierarchies. This illustrates that non-neural systems can behave as if they obeyed modular principles. In this section, we attempt to connect this back to the role of factorised dynamics in nervous systems.

The first point of contact is the role of local, reciprocal, interactions [59] as seen in the density dynamics. In this setting, a mean-field is essentially a description of the message passed to a given neuronal population. In a dynamical formulation, the gradients of the Hamiltonian potentials that comprise this mean-field are interpretable as synaptic weights. Figure 6 unpacks a neuronal network whose dynamics correspond to those above. The graphic on the left shows the interaction between expectations of a single factor and one of its neighbours that shares a local potential (i.e., a constituent of its Markov blanket [60]—the set of states that insulates a node from the rest of the network). Here, each factor may be thought of as predicting the other (via the η functions). This prediction is subtracted from the current expectation (μ) to give an error term (ε). The assumption here is that the time constants of the neural populations representing this error are very short relative to those of the expectations. The error term induces updates in the expectation such that it conforms to the prediction. This is a very simple (linear) form of predictive coding [61,62]—a prominent theory of brain function.

The central image shows what happens when there are multiple constituents to each Markov blanket. There are two ways in which this may manifest anatomically. The first, shown in each of the sensory streams, is that the error populations may accumulate predictions from each blanket constituent. The second is shown in the centre, where multisensory integration takes place. Here, there are multiple error terms, one from each constituent of the blanket. Intuitively, this is as if each error term gets a vote on the expectation, and the resulting attracting point is some combination of these. The influence of the error neurons on those populations representing expectations inherits the solenoidal and diffusion tensor terms. These may be interpreted as intrinsic (within-region) connectivity. The influence of these is shown in the graphic on the right of Figure 6. When the amplitude of fluctuations is large, the error neurons drive the expectations rapidly to their fixed point. However, when the solenoidal term is large, the reciprocal excitatory–inhibitory loop dominates, promoting oscillatory activity. Together, these terms contribute the damped oscillations that underwrite evoked response potentials in electrophysiological studies [63].

A second point of contact is that we have focused on the dynamics of conditional densities. The relevance of this is twofold. First, brain dynamics are generally studied by looking at neural responses to sensory stimuli (i.e., the neural dynamics conditioned upon sensation). Second, conditional densities of this kind underwrite the Bayesian brain hypothesis [64–67]. This view suggests that the brain employs a generative model comprising prior and likelihood densities to predict sensory data. This generative model is the Hamiltonian we have been discussing.

Neural dynamics are then interpretable as forming posterior beliefs (conditioned upon data) about the causes of these data. In saying this, we have deliberately conflated two different perspectives on the Bayesian brain: we have interpreted our stochastic system as if it were a nervous system or a neural network. As such, the density dynamics reflect our beliefs, as observers, about that system, not the network’s beliefs about the outside world. In other words, the posterior is the probability of a neural state given an observation y . The other perspective is that, if we interpret the Hamiltonian as a generative model for y , the density dynamics acquire an interpretation as the brain’s inference about the causes of its sensory data.

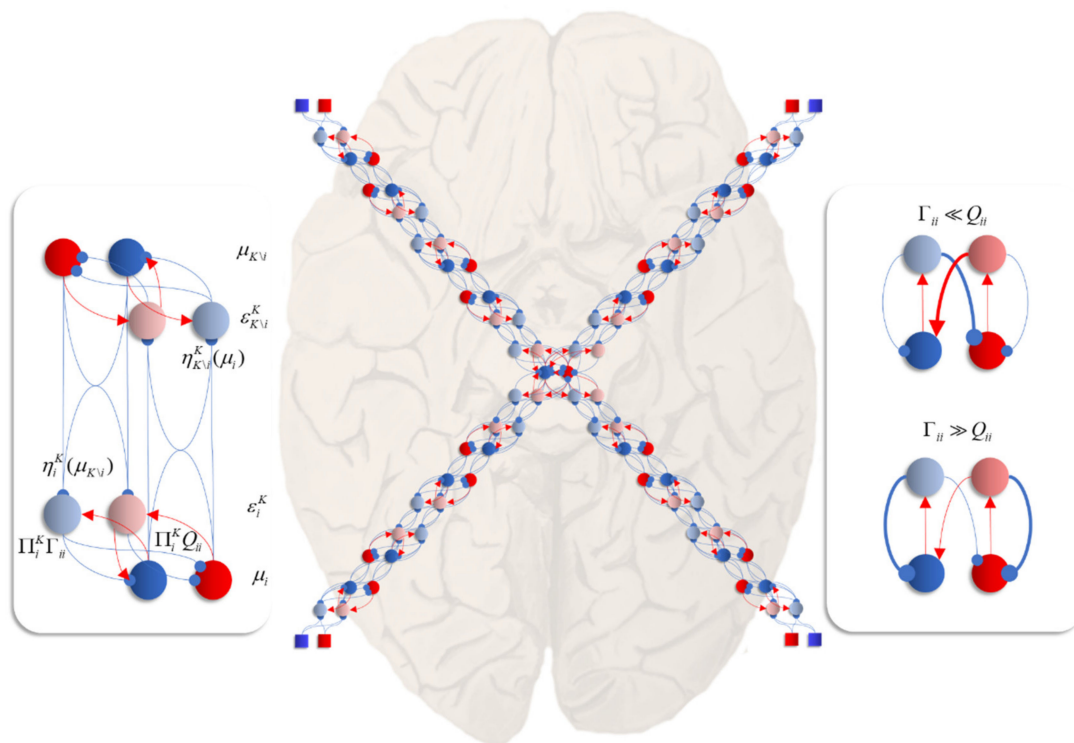


Figure 6. The central panel in this figure shows an interpretation of Equations (13) and (14), applied to the Hamiltonian of Figure 1, as a neuronal network. This shows a reciprocal message passing in which more central and more peripheral regions communicate along a neural hierarchy. Each arm of this hierarchy connects central regions to sensory input (shown as squares, consistent with previous figures). Central regions therefore have multimodal properties, responding to any of the sensory perturbations. Peripheral regions are more specialised in virtue of their proximity to external input. The panel on the left unpacks the connections between two regions (modules, or factors of a mean-field density) in detail. This includes neural populations representing the (2-dimensional) mode (in red and blue), auxiliary variables (in lighter shades) playing the role of prediction errors (i.e., gradients of the local Hamiltonian), and connections between these. Blue connections are inhibitory while red are connections excitatory. Note that, while some populations are shown as giving rise to both excitatory and inhibitory connections, we do not intend to imply a violation of Dale’s law. The assumption here is that there are intermediate inhibitory neurons that act to change the sign of the connection. The panel on the right highlights the importance of intrinsic (intra-modular) connectivity, and the role of the diffusion tensor (Γ) and solenoidal flow (Q) in determining neural activity. If the solenoidal component is large relative to the diffusion tensor, this leads to net excitation of the blue neuron by the red, and net inhibition of the red by the blue. This pattern of connectivity favours intrinsically driven oscillations. The circuit dominated by the diffusion tensor favours rapid convergence of neural activity to a fixed point.

For the Hamiltonian used here, this implies some variable that has consequences for four different sensory modalities (y_1, \dots, y_4). For instance, the position of a cup of coffee has potential consequences for vision, gustation, olfaction, and somatosensation. It may be that the data-generating process is of a form that requires some transformation of the x variables, or even that the generative model is not an accurate description of the data-generating process [68]. Regardless of whether the model is a ‘good’ model, the inferential interpretation is useful in thinking about modularity. This is because it allows us to conceptualise a factor of the system as performing computations *about* something. If each factor is about something different, each can be thought of as a specialised module with a definitive role, in relation to the external environment.

In summary, we can interpret the dynamics of a system described by mean-field density dynamics in terms of messages (i.e., mean-fields) passed between module-like regions of a network [69–71]. For sufficiently sparse conditional dependency structures—like that of the Hamiltonian employed here—the message passing is evocative of synaptic communication in sparse neuronal networks. Interpreted as such, extrinsic (between-node) connection weights are determined by those terms in the Hamiltonian that contribute to a given mean-field. This is distinct to the intrinsic (within-node) connectivity. Intrinsic connections serve to optimise local potentials (given by summing the local mean-fields) through a combination of dissipative (gradient descent) and conservative (solenoidal) flows. Together these ensure a damped oscillation results during return to steady state following a perturbation. Finally, we highlighted the consistency of this perspective with Bayesian theories of brain function, interpreting conditional densities as posterior inferences about the causes of sensory data.

6. Discussion

The key message of this paper is that the concept of a ‘module’ simply refers to a factor of a probability distribution describing a system, and, implicitly, Bayesian beliefs held by a system. To underwrite this argument, we appealed to mean-field theory—a branch of statistical physics that deals with factorisation of probabilistic systems. We illustrated, using a system described by an arbitrarily constructed Hamiltonian, that the density dynamics of a high-dimensional stochastic system may be decomposed into factorised densities of low dimensional components that communicate with one another via their mean-fields. Finally, we interpreted this local communication in terms of synaptic message passing, highlighting the emergent distinction between intrinsic and extrinsic connectivity and the Bayesian interpretation of these dynamics. Crucially, this dynamical and inference architecture depends only on factorisation.

In the above, we have largely ignored the processes generating the variable y , which played the role of sensory data in the final section. While not necessary for the points we sought to address, including these processes has an important consequence for the way in which we think about the dynamics of sentient systems. Specifically, associating average flows of a system, subject to sensory perturbations, with average flows of the data-generating processes enables a reformulation of neuronal message passing in terms of the Hamiltonians of external dynamical systems. The Hamiltonian then becomes synonymous with a generative model of the data generating process. This Bayesian mechanical formulation [16] can be supplemented with the reciprocal influence, to account for neuronal influence on the external world (i.e., action). Things become even more interesting when we think about distributions over alternative trajectories of the internal, active, sensory, and external components of the system [72]. These give rise to the appearance of goal directed and exploratory behaviour. For introductions to the resulting active inference schemes, see [73,74].

We have kept things deliberately simple in the above, through use of a quadratic Hamiltonian. The treatment above, and in particular, the use of a Laplace assumption, retains validity in non-quadratic settings (e.g., [75,76]), but only in regions near the mode of the Hamiltonian. Clearly the assumption of a Gaussian variational density is inappropriate when the system tends towards multimodal densities. This is not a problem for the general principle of factorisation but does mean that solutions based upon the specific formulation of density dynamics used here are only locally valid. For a more general formulation, we could appeal to a more flexible family of variational distributions. An example would be a mixture of Gaussians (of the sort used in clustering applications). These allow for multimodal densities, through a linear combination of Gaussian densities with different modes [77]. In the setting of computational neuroscience, approaches of this sort have been employed to combine models of discrete decision making with those used to solve continuous inference problems [26]. Generative models of this sort have been used to simulate the interface between the selection and enaction of oculomotor saccades and [78], including the performance of oculomotor delay-period tasks [79] like those used in the study of working memory [80–82]. Such mixed models have also been used in the context of modelling neuroimaging data, to understand the way in which the brain switches between

alternative connectivity states [83]. The implication here is that a more comprehensive understanding of the interaction between different factors of a neural system may require some factors representing Gaussian densities, and others categorical distributions over discrete variables.

Another interesting direction in which the formulations above may be extended is in tree decomposition of the Hamiltonians. This addresses the question of how certain kinds of mean-field assumptions (or more sophisticated approximations) may be justified by considering the structure of the Hamiltonian. An important idea here is that of tree-weighted re-parameterisation [84]. This is a class of methods designed to find alternative groupings (i.e., factorisations) of the variables in the graph describing the Hamiltonian. The idea is to create a simpler graph from the original by grouping together highly connected regions of the graph, while allowing for overlaps between groups. These methods provide an alternative perspective on the variational distributions in Table 2. Each Kikuchi approximation may be thought of as an alternative tree-weighted re-parametrisation. The utility of this perspective is that choices of alternative variational densities or trees may be scored. This scoring ends up approximating a KL-Divergence between the distributions under different parameterisations of the tree [85]—sharing the same fixed points as the associated free energy. As such, these techniques could be used to find the ‘best’ decomposition of a system. Another perspective on the same problem is that this decomposition rests upon finding a decomposition based upon Markov blankets in a dynamic setting. This uses adjacency matrices based upon a system Jacobian to find a decomposition such that each blanketed structure in the network is independent of all other structures given their blanket. For a numerical proof-of-principle of this approach, see [86].

Finally, it is worth considering what is gained by thinking about brain function in terms of interacting factors of a probability density. Ultimately, the gains are very similar to that of modularisation. From a neuroscientific point of view, to understand connectivity in the brain, it is necessary to know what the things being connected are [2]. In addition, it is useful to know that some aspects of brain function may be usefully studied in isolation, before placing this in the context of the wider neuronal network. More broadly, factorisation underwrites the notion of transfer learning or context invariance [87,88]. This is the idea that knowledge acquired in one context may be transferred over to a new scenario. Put simply, if we learn that water boils at a temperature of around 100 °C, it should not matter if we change the context by moving to a new location. In the absence of transfer learning, this would have to be learned again in the new context. Each combination of location and temperature would be associated with its own belief about the likelihood of water having boiled. However, simply by factorising temperature and location, we can transfer our beliefs about the relationship between temperature and boiling to any location, c.f., carving nature at its joints via factorisation. Of course, moving to a location at a different altitude does change the temperature at which water boils. This is where the mean-field communication between factors becomes important, correcting for the drastic commitment to think of location and temperature as independent variables. While a trivial example, this highlights the fundamental relationship between factorisation and domain generality. The advantage of framing these problems explicitly in terms of mean-field theory, as opposed to modularity, is that it comes along with a well-established mathematical framework, whose legacy can be traced back to Occam’s principle [89] and the maximum entropy principle [90]. The simplicity of this perspective rests upon Equation (4) and the notion of factorisation. Both modular and mean-field accounts implicitly appeal to factorisations that enable descriptions of parts of a system (modules or marginals). The mean-field perspective is attractive because it does not require additional assumptions. It reformulates the challenge of understanding brain function to one of specifying the Hamiltonian (generative model) that the brain must solve and the variational distribution most appropriate for doing so. This sidesteps the anthropomorphised and *ad hoc* nature of modular accounts, in favour of a formalism grounded in the statistical physics of self-evidencing [91].

7. Conclusions

While not a definitive rejection of a modular perspective on brain function, the treatment presented here suggest that a simpler framing is in terms of factorisation and communication via mean-fields. The mean-field formulation preserves the notions of modular specialisation and information encapsulation. It allows us to work with probability densities within a factor of the variational distribution but does not require propagation of the full density between factors. This ensures a low dimensional passing of messages between factors, just as modules are thought to summarise the output of internal computations for the benefit of their neighbours. This provides a point of connection between Bayesian theories of brain function and the statistical message passing schemes thought to underwrite synaptic communication and computation. In short, the modular view of brain function may be the result of an intuitive application of mean-field theory. In making this explicit, we can draw upon developments in stochastic physics and develop a more formal, quantitative, account of neuronal organisation, from first principles.

Author Contributions: Conceptualization, T.P., N.S., and K.J.F; Formal analysis, T.P.; Software, T.P.; Writing—original draft, T.P.; Writing—review & editing, N.S. and K.J.F. All authors have read and agreed to the published version of the manuscript.

Funding: KJF is a Wellcome Principal Research Fellow (Ref: 088130/Z/09/Z). NS is funded by the Medical Research Council (Ref: MR/S502522/1).

Conflicts of Interest: The authors declare no conflict of interest.

Software Note: The simulations presented here may be reproduced and customised from Matlab (R2019a) code available at <https://github.com/tejparr/Modules-or-Mean-Fields>. Figures 2–5 are generated by the DEMO_MeanFieldsModules.m script.

Appendix A

This appendix provides a derivation of the Fokker–Planck Equation that lets us re-express the behaviour of a stochastic system in terms of its (deterministic) density dynamics. This is based upon the (more rigorous) treatment in [30] and is designed to provide some intuition as to the relationship between stochastic differential equations and their density dynamics. First, we note that the probability of x at time τ can be obtained by marginalising a joint density that includes this time and a previous time:

$$p(x, \tau) = \int p(x, \tau | x - \Delta x, \tau - \Delta \tau) p(x - \Delta x, \tau - \Delta \tau) d\Delta x \quad (\text{A1})$$

For the purposes of this appendix, we use the notation $p(x, \tau)$ to mean the probability density of x at time τ . We omit the conditioning on y that appears throughout the main text. Performing a Taylor series expansion (of $z = x - \Delta x$ around x) of the integrand, we can re-write this as:

$$\begin{aligned} p(x, \tau) &= \int \sum_{n=0} \frac{1}{n!} (-\Delta x)^n \nabla_z^n p(z + \Delta x, \tau | z, \tau - \Delta \tau) p(z, \tau - \Delta \tau) \Big|_{z=x} d\Delta x \\ &= \sum_{n=0} \frac{(-1)^n}{n!} \nabla_z^n \int \Delta x^n p(z + \Delta x, \tau | z, \tau - \Delta \tau) d\Delta x p(z, \tau - \Delta \tau) \Big|_{z=x} \\ &= \sum_{n=0} \frac{(-1)^n}{n!} \nabla_z^n \mathbb{E}_{p(z+\Delta x, \tau | z, \tau - \Delta \tau)} [\Delta x^n] p(z, \tau - \Delta \tau) \Big|_{z=x} \\ &= \sum_{n=0} \frac{(-1)^n}{n!} \nabla_x^n \mathbb{E}_{p(\Delta x | \Delta \tau)} [\Delta x^n] p(x, \tau - \Delta \tau) \end{aligned} \quad (\text{A2})$$

Subtracting the first term of the sum from both sides, we get:

$$p(x, \tau) - p(x, \tau - \Delta \tau) = \sum_{n=1} \frac{(-1)^n}{n!} \nabla_x^n \mathbb{E}_{p(\Delta x | \Delta \tau)} [\Delta x^n] p(x, \tau - \Delta \tau) \quad (\text{A3})$$

From here, we can find the form of the rate of change of the probability density by taking limits:

$$\begin{aligned} \dot{p}(x, \tau) &= \lim_{\Delta \tau \rightarrow 0} \left\{ \frac{1}{\Delta \tau} (p(x, \tau) - p(x, \tau - \Delta \tau)) \right\} \\ &= \sum_{n=1} \frac{(-1)^n}{n!} \nabla_x^n \lim_{\Delta \tau \rightarrow 0} \left\{ \mathbb{E}_{p(\Delta x | \Delta \tau)} \left[\frac{\Delta x^n}{\Delta \tau} \right] \right\} p(x, \tau) \end{aligned} \quad (\text{A4})$$

For the first two terms in the expansion, we have:

$$\begin{aligned}
 \lim_{\Delta\tau \rightarrow 0} \left\{ \mathbb{E}_{p(\Delta x|\Delta\tau)} \left[\frac{\Delta x}{\Delta\tau} \right] \right\} &= \lim_{\Delta\tau \rightarrow 0} \left\{ \mathbb{E}_{p(\Delta x|\Delta\tau)} \left[\frac{\int_{\tau}^{\tau+\Delta\tau} f(x,y) + \omega(t) dt}{\Delta\tau} \right] \right\} = f(x,y) \\
 \lim_{\Delta\tau \rightarrow 0} \left\{ \mathbb{E}_{p(\Delta x|\Delta\tau)} \left[\frac{\Delta x^2}{\Delta\tau} \right] \right\} &= \lim_{\Delta\tau \rightarrow 0} \left\{ \underbrace{\mathbb{E}_{p(\Delta x|\Delta\tau)} \left[\frac{\int_{\tau}^{\tau+\Delta\tau} f(x,y) dt \int_{\tau}^{\tau+\Delta\tau} f(x,y) dt}{\Delta\tau} \right]}_{=0} \right. \\
 &\quad \left. + 2 \lim_{\Delta\tau \rightarrow 0} \left\{ \mathbb{E}_{p(\Delta x|\Delta\tau)} \left[\frac{\int_{\tau}^{\tau+\Delta\tau} f(x,y) dt \int_{\tau}^{\tau+\Delta\tau} \omega(t) dt}{\Delta\tau} \right] \right\} \right. \\
 &\quad \left. + \lim_{\Delta\tau \rightarrow 0} \left\{ \mathbb{E}_{p(\Delta x|\Delta\tau)} \left[\frac{\int_{\tau}^{\tau+\Delta\tau} \omega(t) dt \int_{\tau}^{\tau+\Delta\tau} \omega(t) dt}{\Delta\tau} \right] \right\} \right\} \\
 &= \lim_{\Delta\tau \rightarrow 0} \left\{ \mathbb{E}_{p(\Delta x|\Delta\tau)} \left[\frac{\int_{\tau}^{\tau+\Delta\tau} \int_{\tau}^{\tau+\Delta\tau} \omega(t)\omega(s) dt ds}{\Delta\tau} \right] \right\} \\
 &= \lim_{\Delta\tau \rightarrow 0} \left\{ \frac{\int_{\tau}^{\tau+\Delta\tau} \int_{\tau}^{\tau+\Delta\tau} 2\Gamma\delta(t-s) dt ds}{\Delta\tau} \right\} = 2\Gamma
 \end{aligned} \tag{A5}$$

As such, the density dynamics (up to the second order expansion) may be expressed as follows:

$$\dot{p}(x, \tau) = \nabla \cdot (\Gamma \nabla - f(x)) p(x, \tau) \tag{A6}$$

This is the Fokker–Planck equation. Its utility is that, in place of studying specific instances of a stochastic trajectory, we can work with a deterministic equation that tells us how densities change over time. It may seem arbitrary to truncate the expansion in Equation A4 after the second term. The reason for doing so is that, by the Pawula theorem, additional terms support evolution to densities that are inconsistent with probability densities unless an infinite number of terms are included to preclude this.

Appendix B

To use a Fokker–Planck formulation of dynamics practically, it is often necessary to find some parameterisation of the probability density such that the rate of change of the density may be reformulated in terms of the rate of change of the parameters of that density. One of the simplest options here is to take a Taylor series approximation to the log probability density. When this is truncated after the quadratic term, this is known as a Laplace approximation [38]. Here, we assume the log variational density is quadratic:

$$\begin{aligned}
 \ln q(x_i|y) &= \ln q(\mu_i|y) + \underbrace{(x_i - \mu_i) \cdot \nabla_{x_i} \ln q(x_i|y)}_{=0} \Big|_{x_i=\mu_i} - \frac{1}{2} (x_i - \mu_i) \cdot \Sigma_{ii}^{-1} (x_i - \mu_i) \\
 \Sigma_{ii}^{-1} &\triangleq -\nabla_{x_i x_i} \ln q(x_i|y) \Big|_{x_i=\mu_i} \\
 \mu_i &\triangleq \operatorname{argmax}_{x_i} \{q(x_i|y)\} \\
 \Rightarrow q(x_i|y) &= \mathcal{N}(x_i; \mu_i, \Sigma_i^{-1})
 \end{aligned} \tag{A7}$$

When dealing with linear dynamical systems, the Laplace approximation is exact. More generally, it is suitable for describing systems in the vicinity of the mode of the variational density. With this

assumption in place, we can find expressions for the rate of change of the sufficient statistics of the probability density:

$$\begin{aligned} \mu_i &\triangleq E_{q(x_i|y)}[x_i] = \int_{-\infty}^{\infty} x_i q(x_i|y) dx_i \\ \dot{\mu}_i &= \int_{-\infty}^{\infty} x_i \dot{q}(x_i|y) dx_i \\ &= \int_{-\infty}^{\infty} x_i \nabla_{x_i} \cdot (\Gamma_{ii} \nabla_{x_i} q(x_i|y)) dx_i + \beta \int_{-\infty}^{\infty} x_i \nabla_{x_i} \cdot ((\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i, y) q(x_i|y)) dx_i \end{aligned} \tag{A8}$$

Integration by parts gives:

$$\begin{aligned} \dot{\mu}_i &= \underbrace{\left[x_i \Gamma_{ii} \nabla_{x_i} q(x_i|y) \right]_{-\infty}^{\infty}}_{=0} - \underbrace{\int_{-\infty}^{\infty} \Gamma_{ii} \nabla_{x_i} q(x_i|y) dx_i}_{=0} \\ &\quad \underbrace{\beta \left[x_i (\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i, y) q(x_i|y) \right]_{-\infty}^{\infty}}_{=0} - \beta \int_{-\infty}^{\infty} (\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i, y) q(x_i|y) dx \\ &= -\beta (\Gamma_{ii} - Q_{ii}) E_{q(x_i|y)} \left[\nabla_{x_i} h_i(x_i, y) \right] \end{aligned} \tag{A9}$$

The same procedure can then be applied to the covariance:

$$\begin{aligned} \Sigma_{ii} &\triangleq E_{q(x_i|y)}[\Delta x_i \Delta x_i^T] = \int_{-\infty}^{\infty} \Delta x_i \Delta x_i^T q(x_i|y) dx_i \\ \dot{\Sigma}_{ii} &= \int_{-\infty}^{\infty} \Delta x_i \Delta x_i^T \dot{q}(x_i|y) dx_i \\ &= \int_{-\infty}^{\infty} \Delta x_i \Delta x_i^T \nabla_{x_i} \cdot (\Gamma_{ii} \nabla_{x_i} q(x_i|y)) dx_i \\ &\quad + \beta \int_{-\infty}^{\infty} \Delta x_i \Delta x_i^T \nabla_{x_i} \cdot ((\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i, y) q(x_i|y)) dx_i \end{aligned} \tag{A10}$$

Again, integrating by parts gives:

$$\begin{aligned} \dot{\Sigma}_{ii} &= \underbrace{\left[\Delta x_i \Delta x_i^T \Gamma_{ii} \nabla_{x_i} q(x_i|y) \right]_{-\infty}^{\infty}}_{=0} - \int_{-\infty}^{\infty} \left(\Delta x_i (\Gamma_{ii} \nabla_{x_i} q(x_i|y))^T + \Gamma_{ii} \nabla_{x_i} q(x_i|y) \Delta x_i^T \right) dx_i \\ &\quad + \beta \underbrace{\left[\Delta x_i \Delta x_i^T (\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i, y) q(x_i|y) \right]_{-\infty}^{\infty}}_{=0} \\ &\quad - \beta \int_{-\infty}^{\infty} \left(\Delta x_i ((\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i, y))^T + (\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i, y) \Delta x_i^T \right) q(x_i|y) dx_i \\ &= -2\Gamma_{ii} \underbrace{\left[\Delta x_i q(x_i|y) \right]_{-\infty}^{\infty}}_{=0} + 2\Gamma_{ii} \underbrace{\int_{-\infty}^{\infty} q(x_i|y) dx_i}_{=1} \\ &\quad - \beta E_{q(x_i|y)} \left[\left(\Delta x_i ((\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i, y))^T + (\Gamma_{ii} - Q_{ii}) \nabla_{x_i} h_i(x_i, y) \Delta x_i^T \right) \right] \\ &= 2\Gamma_{ii} \\ &\quad - \beta E_{q(x_i|y)} \left[\Delta x_i \nabla_{x_i} h_i(x_i, y)^T \right] (\Gamma_{ii} - Q_{ii})^T \\ &\quad - \beta (\Gamma_{ii} - Q_{ii}) E_{q(x_i|y)} \left[\nabla_{x_i} h_i(x_i, y) \Delta x_i^T \right] \end{aligned} \tag{A11}$$

Together, Equations (A9) and (A11) provide parameterised forms for the density dynamics, and a tractable method of numerically integrating a Fokker–Planck Equation.

References

1. Fodor, J.A. *The Modularity of Mind: An Essay on Faculty Psychology, Reprint ed.*; MIT Press: Cambridge, MA, USA, 1983.
2. Friston, K.J.; Price, C.J. Modules and brain mapping. *Cogn. Neuropsychol.* **2011**, *28*, 241–250. [[CrossRef](#)] [[PubMed](#)]
3. Clune, J.; Mouret, J.-B.; Lipson, H. The evolutionary origins of modularity. *Biol. Sci.* **2013**, *280*, 20122863. [[CrossRef](#)] [[PubMed](#)]
4. Hipolito, I.; Kirchhoff, M.D. The Predictive Brain: A Modular View of Brain and Cognitive Function? preprints, 2019. Available online: <https://www.preprints.org/manuscript/201911.0111/v1> (accessed on 13 May 2020).
5. Baltieri, M.; Buckley, C.L. The modularity of action and perception revisited using control theory and active inference. In *Artificial Life Conference Proceedings*; MIT Press: Cambridge, MA, USA, 2018; pp. 121–128.
6. Cosmides, L.; Tooby, J. Origins of domain specificity: The evolution of functional organization. In *Mapping the Mind: Domain Specificity in Cognition and Culture*; Cambridge University Press: New York, NY, USA, 1994; pp. 85–116.
7. Weiss, P. L'hypothèse du champ moléculaire et la propriété ferromagnétique. *J. Phys. Theor. Appl.* **1907**, *6*, 661–690. [[CrossRef](#)]
8. Kadanoff, L.P. More is the Same; Phase Transitions and Mean Field Theories. *J. Stat. Phys.* **2009**, *137*, 777. [[CrossRef](#)]
9. Cessac, B. Mean Field Methods in Neuroscience. 2015. Available online: <https://core.ac.uk/download/pdf/52775181.pdf> (accessed on 13 May 2020).
10. Fasoli, D. Attacking the Brain with Neuroscience: Mean-Field Theory, Finite Size Effects and Encoding Capability of Stochastic Neural Networks. Ph.D. Thesis, Université Nice Sophia Antipolis, Nice, France, 2013.
11. Winn, J.; Bishop, C.M. Variational message passing. *J. Mach. Learn. Res.* **2005**, *6*, 661–694.
12. Gadomski, A.; Kruszewska, N.; Ausloos, M.; Tadych, J. On the Harmonic-Mean Property of Model Dispersive Systems Emerging Under Mononuclear, Mixed and Polynuclear Path Conditions. In *Traffic and Granular Flow'05*; Springer: Berlin/Heidelberg, Germany, 2007.
13. Hethcote, H.W. Three Basic Epidemiological Models. In *Applied Mathematical Ecology*; Levin, S.A., Hallam, T.G., Gross, L.J., Eds.; Springer: Berlin/Heidelberg, Germany, 1989; pp. 119–144.
14. Lasry, J.-M.; Lions, P.-L. Mean field games. *Jpn. J. Math.* **2007**, *2*, 229–260. [[CrossRef](#)]
15. Lelarge, M.; Bolot, J. A local mean field analysis of security investments in networks. In Proceedings of the 3rd international workshop on Economics of networked systems, Seattle, WA, USA, 20–22 August 2008.
16. Friston, K. A free energy principle for a particular physics. *arXiv* **2019**, arXiv:1906.10184.
17. Yoshioka, D. The Partition Function and the Free Energy. In *Statistical Physics: An Introduction*; Yoshioka, D., Ed.; Springer: Berlin/Heidelberg, Germany, 2007; pp. 35–44.
18. Hinton, G.E.; Zemel, R.S. Autoencoders, minimum description length and Helmholtz free energy. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 1994.
19. Beal, M.J. *Variational Algorithms for Approximate Bayesian Inference*; University of London: London, UK, 2003.
20. Bogolyubov, N.N. On model dynamical systems in statistical mechanics. *Physica* **1966**, *32*, 933–944. [[CrossRef](#)]
21. Feynman, R.P. Space-Time Approach to Non-Relativistic Quantum Mechanics. *Rev. Mod. Phys.* **1948**, *20*, 367–387. [[CrossRef](#)]
22. Loeliger, H. An introduction to factor graphs. *IEEE Signal Process. Mag.* **2004**, *21*, 28–41. [[CrossRef](#)]
23. Vontobel, P.O. A factor-graph approach to Lagrangian and Hamiltonian dynamics. In *2011 IEEE International Symposium on Information Theory Proceedings*; IEEE: Piscataway, NJ, USA, 2011.
24. Loeliger, H.; Vontobel, P.O. Factor Graphs for Quantum Probabilities. *IEEE Trans. Inf. Theory* **2017**, *63*, 5642–5665. [[CrossRef](#)]
25. Parr, T.; Friston, K.J. The Anatomy of Inference: Generative Models and Brain Structure. *Front. Comput. Neurosci.* **2018**, *12*, 90. [[CrossRef](#)] [[PubMed](#)]
26. Friston, K.J.; Parr, T.; de Vries, B. The graphical brain: Belief propagation and active inference. *Netw. Neurosci.* **2017**, *1*, 381–414. [[CrossRef](#)] [[PubMed](#)]
27. Pelizzola, A. Cluster variation method in statistical physics and probabilistic graphical models. *J. Phys. A Math. Gen.* **2005**, *38*, R309–R339. [[CrossRef](#)]

28. Yedidia, J.S.; Freeman, W.T.; Weiss, Y. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Trans. Inf. Theory* **2005**, *51*, 2282–2312. [[CrossRef](#)]
29. Frey, B.J.; MacKay, D.J.C. A revolution: Belief propagation in graphs with cycles. In Proceedings of the 1997 Conference on Advances in Neural Information Processing Systems 10; MIT Press: Denver, CL, USA, 1998; pp. 479–485.
30. Risken, H. Fokker-Planck Equation. In *The Fokker-Planck Equation: Methods of Solution and Applications*; Springer: Berlin/Heidelberg, Germany, 1996; pp. 63–95.
31. Ao, P. Potential in stochastic differential equations: Novel construction. *J. Phys. A Math. Gen.* **2004**, *3*, L25–L30. [[CrossRef](#)]
32. Kwon, C.; Ao, P.; Thouless, D.J. Structure of stochastic dynamics near fixed points. *Proc. Natl. Acad. Sci. USA* **2005**, *102*, 13029–13033. [[CrossRef](#)]
33. Ma, Y.-A.; Chen, T.; Fox, E. A complete recipe for stochastic gradient MCMC. In *Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2015.
34. Pylyshyn, Z. Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behav. Brain Sci.* **1999**, *22*, 341–365. [[CrossRef](#)]
35. Seifert, U. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.* **2012**, *75*, 126001. [[CrossRef](#)]
36. Grzelczak, M.; Vermant, J.; Furst, E.M.; Liz-Marzán, L.M. Directed Self-Assembly of Nanoparticles. *ACS Nano* **2010**, *4*, 3591–3605. [[CrossRef](#)]
37. Cheng, J.Y.; Mayes, A.M.; Ross, C.A. Nanostructure engineering by templated self-assembly of block copolymers. *Nat. Mater.* **2004**, *3*, 823–828. [[CrossRef](#)] [[PubMed](#)]
38. Marreiros, A.C.; Kiebel, S.J.; Daunizeau, J.; Harrison, L.M.; Friston, K.J. Population dynamics under the Laplace assumption. *Neuroimage* **2009**, *44*, 701–714. [[CrossRef](#)] [[PubMed](#)]
39. Moran, R.; Pinotsis, D.A.; Friston, K. Neural masses and fields in dynamic causal modeling. *Front. Comput. Neurosci.* **2013**, *7*, 57. [[CrossRef](#)] [[PubMed](#)]
40. Hastings, W.K. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **1970**, *57*, 97–109. [[CrossRef](#)]
41. Yildirim, I. *Bayesian Inference: Gibbs Sampling*; Technical Note; University of Rochester: Rochester, NY, USA, 2012.
42. Neal, R.M. *Probabilistic Inference Using Markov Chain Monte Carlo Methods*; Department of Computer Science, University of Toronto: Toronto, ON, Canada, 1993.
43. Girolami, M.; Calderhead, B. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Stat. Soc. Ser. B* **2011**, *73*, 123–214. [[CrossRef](#)]
44. Ungerleider, L.G.; Haxby, J.V. ‘What’ and ‘where’ in the human brain. *Curr. Opin. Neurobiol.* **1994**, *4*, 157–165. [[CrossRef](#)]
45. Winkler, I.; Denham, S.; Mill, R.; Böhm, T.M.; Bendixen, A. Multistability in auditory stream segregation: A predictive coding view. *Philos. Trans. R. Soc. B Biol. Sci.* **2012**, *367*, 1001–1012. [[CrossRef](#)]
46. Hickok, G.; Poeppel, D. Dorsal and ventral streams: A framework for understanding aspects of the functional anatomy of language. *Cognition* **2004**, *92*, 67–99. [[CrossRef](#)]
47. Friston, K.; Buzsáki, G. The Functional Anatomy of Time: What and When in the Brain. *Trends Cogn. Sci.* **2016**, *20*, 500–511. [[CrossRef](#)]
48. Kiebel, S.J.; Daunizeau, J.; Friston, K.J. A Hierarchy of Time-Scales and the Brain. *PLoS Comput. Biol.* **2008**, *4*, e1000209. [[CrossRef](#)]
49. Cocchi, L.; Sale, M.V.; Gollo, L.L.; Bell, P.T.; Nguyen, V.T.; Zalesky, A.; Breakspear, M.; Mattingley, J.B. A hierarchy of timescales explains distinct effects of local inhibition of primary visual cortex and frontal eye fields. *eLife* **2016**, *5*, e15252. [[CrossRef](#)] [[PubMed](#)]
50. Hasson, U.; Yang, E.; Vallines, I.; Heeger, D.J.; Rubin, N. A Hierarchy of Temporal Receptive Windows in Human Cortex. *Off. J. Soc. Neurosci.* **2008**, *28*, 2539–2550. [[CrossRef](#)] [[PubMed](#)]
51. Murray, J.D.; Bernacchia, A.; Freedman, D.J.; Romo, R.; Wallis, J.D.; Cai, X.; Padoa-Schioppa, C.; Pasternak, T.; Seo, H.; Lee, D.; et al. A hierarchy of intrinsic timescales across primate cortex. *Nat. Neurosci.* **2014**, *17*, 1661–1663. [[CrossRef](#)] [[PubMed](#)]
52. Murata, A.; Fadiga, L.; Fogassi, L.; Gallese, V.; Raos, V.; Rizzolatti, G. Object representation in the ventral premotor cortex (area F5) of the monkey. *J. Neurophysiol.* **1997**, *78*, 2226–2230. [[CrossRef](#)] [[PubMed](#)]

53. Giard, M.H.; Peronnet, F. Auditory-Visual Integration during Multimodal Object Recognition in Humans: A Behavioral and Electrophysiological Study. *J. Neurophysiol.* **1999**, *11*, 473–490. [[CrossRef](#)] [[PubMed](#)]
54. Wallace, M.T.; Meredith, M.A.; Stein, B.E. Multisensory Integration in the Superior Colliculus of the Alert Cat. *J. Neurophysiol.* **1998**, *80*, 1006–1010. [[CrossRef](#)]
55. Limanowski, J.; Blankenburg, F. Integration of Visual and Proprioceptive Limb Position Information in Human Posterior Parietal, Premotor, and Extrastriate Cortex. *Off. J. Soc. Neurosci.* **2016**, *36*, 2582–2589. [[CrossRef](#)]
56. Stein, B.E.; Stanford, T.R. Multisensory integration: Current issues from the perspective of the single neuron. *Nat. Rev. Neurosci.* **2008**, *9*, 255–266. [[CrossRef](#)]
57. Tononi, G.; Sporns, O.; Edelman, G.M. A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proc. Natl. Acad. Sci. USA* **1994**, *91*, 5033–5037. [[CrossRef](#)]
58. Fukushima, M.; Betzel, R.F.; He, Y.; van den Heuvel, M.P.; Zuo, X.N.; Sporns, O. Structure-function relationships during segregated and integrated network states of human brain functional connectivity. *Brain Struct. Funct.* **2018**, *223*, 1091–1106. [[CrossRef](#)]
59. Markov, N.T.;ERCsey-Ravasz, M.; Van Essen, D.C.; Knoblauch, K.; Toroczka, Z.; Kennedy, H. Cortical high-density counterstream architectures. *Science* **2013**, *342*, 1238406. [[CrossRef](#)] [[PubMed](#)]
60. Pearl, J. Probabilistic Reasoning. In *Intelligent Systems: Networks of Plausible Inference*; Morgan Kaufmann: San Francisco, CA, USA, 1988.
61. Friston, K.; Kiebel, S. Predictive coding under the free-energy principle. *Philos. Trans. R. Soc. B Biol. Sci.* **2009**, *364*, 1211–1221. [[CrossRef](#)] [[PubMed](#)]
62. Rao, R.P.; Ballard, D.H. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **1999**, *2*, 79–87. [[CrossRef](#)] [[PubMed](#)]
63. David, O.; Kilner, J.M.; Friston, K.J. Mechanisms of evoked and induced responses in MEG/EEG. *NeuroImage* **2006**, *31*, 1580–1591. [[CrossRef](#)] [[PubMed](#)]
64. Knill, D.C.; Pouget, A. The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends Neurosci.* **2004**, *27*, 712–719. [[CrossRef](#)] [[PubMed](#)]
65. Doya, K. *Bayesian Brain: Probabilistic Approaches to Neural Coding*; MIT Press: Cambridge, MA, USA, 2007.
66. Friston, K. The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **2010**, *11*, 127–138. [[CrossRef](#)]
67. O'Reilly, J.X.; Jbabdi, S.; Behrens, T.E.J. How can a Bayesian approach inform neuroscience? *Eur. J. Neurosci.* **2012**, *35*, 1169–1179. [[CrossRef](#)]
68. Tschantz, A.; Seth, A.K.; Buckley, C.L. Learning action-oriented models through active inference. *bioRxiv* **2019**. [[CrossRef](#)]
69. George, D.; Hawkins, J. Towards a mathematical theory of cortical micro-circuits. *PLoS Comput. Biol.* **2009**, *5*, e1000532. [[CrossRef](#)]
70. Parr, T.; Markovic, D.; Kiebel, S.J.; Friston, K.J. Neuronal message passing using Mean-field, Bethe, and Marginal approximations. *Sci. Rep.* **2019**, *9*, 1889. [[CrossRef](#)]
71. Van de Laar, T.W.; de Vries, B. Simulating Active Inference Processes by Message Passing. *Front. Robot. AI* **2019**, *6*, 20. [[CrossRef](#)]
72. Parr, T.; Costa, L.D.; Friston, K. Markov blankets, information geometry and stochastic thermodynamics. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **2020**, *378*, 20190159. [[CrossRef](#)] [[PubMed](#)]
73. Sajid, N.; Ball, P.J.; Friston, K.J. Demystifying active inference. *arXiv* **2019**, arXiv:1909.10863.
74. Da Costa, L.; Parr, T.; Sajid, N.; Veselic, S.; Neacsu, V.; Friston, K. Active inference on discrete state-spaces: A synthesis. *arXiv* **2020**, arXiv:2001.07203.
75. Harding, M.C.; Hausman, J. Using a Laplace: Approximation to Estimate the Random Coefficients logit model by Nonlinear Least Squares*. *Int. Econ. Rev.* **2007**, *48*, 1311–1328. [[CrossRef](#)]
76. Daunizeau, J.; Friston, K.J.; Kiebel, S.J. Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. *Phys. D Nonlinear Phenom.* **2009**, *238*, 2089–2118. [[CrossRef](#)]
77. He, X.; Cai, D.; Shao, Y.; Bao, H.; Han, J. Laplacian regularized gaussian mixture model for data clustering. *IEEE Trans. Knowl. Data Eng.* **2010**, *23*, 1406–1418. [[CrossRef](#)]
78. Parr, T.; Friston, K.J. The Discrete and Continuous Brain: From Decisions to Movement—And Back Again. *Neural Comput.* **2018**, *30*, 2319–2347. [[CrossRef](#)]

79. Parr, T.; Friston, K.J. The computational pharmacology of oculomotion. *Psychopharmacology* **2019**, *236*, 2473–2484. [[CrossRef](#)]
80. Tsujimoto, S.; Postle, B.R. The prefrontal cortex and oculomotor delayed response: A reconsideration of the “mnemonic scotoma”. *J. Cogn. Neurosci.* **2012**, *24*, 627–635. [[CrossRef](#)]
81. Funahashi, S. Functions of delay-period activity in the prefrontal cortex and mnemonic scotomas revisited. *Front. Syst. Neurosci.* **2015**, *9*, 2. [[CrossRef](#)] [[PubMed](#)]
82. Kojima, S.; Goldman-Rakic, P.S. Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response. *Brain Res.* **1982**, *248*, 43–50. [[CrossRef](#)]
83. Zarghami, T.S.; Friston, K.J. Dynamic effective connectivity. *NeuroImage* **2020**, *207*, 116453. [[CrossRef](#)] [[PubMed](#)]
84. Wu, C.-H.; Doerschuk, P.C. Tree approximations to Markov random fields. *IEEE Trans. Pattern Anal. Mach. Intell.* **1995**, *17*, 391–402. [[CrossRef](#)]
85. Wainwright, M.J.; Jaakkola, T.S.; Willsky, A.S. Tree-based reparameterization framework for analysis of sum-product and related algorithms. *IEEE Trans. Inf. Theory* **2003**, *49*, 1120–1146. [[CrossRef](#)]
86. Friston, K. Life as we know it. *J. R. Soc. Interface* **2013**, *10*, 20130475. [[CrossRef](#)]
87. Rojas-Carulla, M.; Schölkopf, B.; Turner, R.; Peters, J. Invariant models for causal transfer learning. *J. Mach. Learn. Res.* **2018**, *19*, 1309–1342.
88. Bengio, Y. Deep learning of representations for unsupervised and transfer learning. *Workshop Conf. Proc.* **2012**, *27*, 17–37.
89. Maisto, D.; Donnarumma, F.; Pezzulo, G. Divide et impera: Subgoalng reduces the complexity of probabilistic inference and problem solving. *J. R. Soc. Interface* **2015**, *12*, 20141335. [[CrossRef](#)]
90. Jaynes, E.T. Information Theory and Statistical Mechanics. *Phys. Rev. Ser. II* **1957**, *106*, 620–630. [[CrossRef](#)]
91. Hohwy, J. The Self-Evidencing Brain. *Noûs* **2016**, *50*, 259–285. [[CrossRef](#)]



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).