

Characterizing Brain Activity at the Mesoscopic Scale

Comparing a Kuramoto model to Resting State MEG data

Daniel Manson

August 2012

Abstract

Despite impressive advances over the last several decades, to a large extent the human brain remains something of a black-box. This state of affairs is easily understood given that the ‘machine inside the box’ consists of a network of approximately 10^{11} nodes, each of which transmits electrical signals in a highly non-linear and ever changing manner. Investigation of individual cell types has been relatively successful in primary sensory areas and hippocampus, but there is yet to be any significant progress in understanding the dynamics of small populations of cells. Appropriate grouping of these small populations results in a brain divided into approximately 1000 regions. These regions can then be further aggregated to produce a labeled map of the brain at the mesoscopic scale where there are about 66 distinct regions. EEG, MEG and fMRI make it possible to record brain activity at this coarse level of network description. In this work we model the 66 regions as a network of weakly coupled oscillators, i.e as a Kuramoto model, and compare model simulations to MEG data collected at resting state.

Contents

1	Introduction	2
2	Models of simple interacting entities	2
2.1	Cellular Automata	2
2.2	Ising Model	2
2.3	Kuramoto Oscillators	3
2.4	Vicsek Particles	4
2.5	Integrate and Fire Neurons	4
3	Mathematical Concepts	5
3.1	Networks	5
3.2	Criticality	6
3.3	Information Theoretic Measures	7
4	Biology of the Brain	9
4.1	Mesoscopic Architecture of the Brain	9
4.2	Brain Waves	10
4.3	The Role of Synchronization	12
5	New Work	13
5.1	Discussion of the model	13
5.2	Comparison with MEG data	14
5.3	Computation and the Kuramoto Model	17
6	Concluding Remarks	19

Abbreviations

MEG	Magnetoencephalography
EEG	Electroencephalography
(f)MRI	(functional) Magnetic Resonance Imaging
DSI	Diffusion Spectrum Imaging
LFP	Local Field Potential
CA	Cellular Automaton
IF	Integrate and Fire (neuron)

1 Introduction

Since the start of the computing era, which can perhaps be dated to the publication in 1937 of Turing’s paper on automatic machines [37], it has been widely understood that a simple set of rules applied on a large enough scale can produce interesting results - a process termed ‘emergence’ [35, p. 14]. And, as Turing himself seemed to believe (also Von Neumann [40]), the human brain is one of the most significant demonstrations of this phenomenon. Thus, in order to understand how the human brain functions it is not enough to catalogue its rule set, we must also find a framework for understanding the larger scale dynamics that the rules produce. These dynamics can be examined at the level of small populations, of perhaps several thousand similar cells. Additionally, they can be examined at a much coarser resolution; at a scale where the brain itself is a network of around 1000 nodes, or even just 66 [15].

In this work we begin by introducing several important models that have the property described above - that is they use simple deterministic rules to produce interesting behavior. From this small collection of models we select one model - the Kuramoto weakly-coupled oscillators - and examine to what extent it helps explain the large scale dynamics of a healthy human brain. Before beginning this examination it will be necessary to introduce a selection of mathematical concepts, namely network theory, criticality, and information theory.

2 Models of simple interacting entities

2.1 Cellular Automata

Introduced in the 1960s, a cellular automaton is an n-dimensional lattice of symbols that evolves in discrete time steps according to a set of rules applied locally at each site in the lattice [43]. The most simple example of this is the 1-dimensional case where there are just two symbols (labeled 0 and 1) and the rules only depend on the cell’s value and its two neighbors’ values, thus there are $2^3 = 8$ rules. Figure 1 shows an example set of rules.

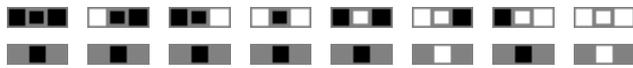


Figure 1: An example set of rules for the binary 1-dimensional cellular automaton. The **top row** shows the cell’s value and the value of its two neighbors, and the **second row** shows the value that the cell should be given at the next time step.

For this type of CA, having only eight rules, with each rule prescribing one of only two possible values there are just $2^8 = 256$ possible rule sets (some of which are equivalent through 0-1 inversion or left-right inversion). Each of these rule sets has been subjected to a significant amount of examination, and famously the rule set known as ‘110’ (see Figure 10 for an example) has been proven to be computationally universal [44]. This means that a large enough lattice evolving according to this rule set can simulate a universal Turing Machine¹ and by extension any algorithm that a modern day computer might run.

CAs of varying dimension, symbol count, and neighborhood size have been used to investigate a wide range of phenomena including, but not limited to, the field of biology [43]. Although in this work we do not compare the brain to CAs we will use CAs as a model system for understanding some mathematical concepts.

2.2 Ising Model

Proposed in the 1920s as a way of studying ferromagnetism, the Ising model is not entirely different from a CA in that it consists of a lattice made up of two symbols that can alternate over the course of time [13]. In the Ising model the two symbols are +1 and -1, corresponding to the magnetic dipole at each point in the material. However, unlike in a CA, the evolution of the system is not governed by simple local rules, instead it depends on the ‘energy’ of the whole system in its current configuration and in potential future configurations, where one of the symbols in the lattice has flipped. Ignoring

¹A Turing Machine [37] is a little like a CA in that it consist of a lattice of symbols that evolve in discrete time steps according to a set of rules. However in the case of a Turing machine the rules are only applied at one special point in the lattice, which acts like a read/write mechanism with a small amount of internal memory. A 1-dimensional system of this kind that has a library of four symbols and seven memory states is known to be universal [43].

external fields, the ‘energy’ of a configuration is defined as the difference between (a) the number of neighboring pairs with the same symbol, and (b) the number of neighboring pairs with different symbols. Transition probabilities also depend on a global parameter that represents temperature.

We will return to the Ising model in a little more detail during the discussion of criticality; here we just give Figure 2 as a brief example of how Ising models behave.

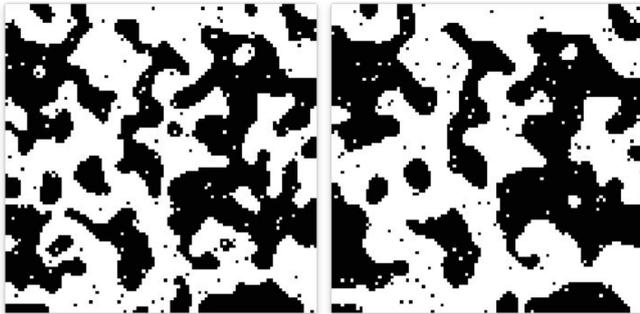


Figure 2: Two frames from a 2-dimensional Ising model simulation (the **left frame** is the initial state, and the **right frame** occurs some short time later). Taken from supplementary video in [4].

2.3 Kuramoto Oscillators

Unlike the two models described above, the Kuramoto model deals with continuous-valued non-spatial entities. Introduced in the 1970s, it was intended to act as a framework for studying synchronous phenomena in the natural world.

At its simplest, the model consists of a number of oscillators, each of which has a phase, θ_i , that traverses the periodic interval $[0, 2\pi]$ at some rate, ω_i (where ω_i is picked from some stated distribution). This in itself is a trivial system; the interesting dynamics only occur when we introduce a weak coupling between the oscillators. The coupling is usually defined as being proportional to the sine of the difference between pairs of oscillators. Thus the model has the form:

$$\frac{d\theta_i}{dt} = \omega_i + K \sum_j \sin(\theta_j - \theta_i) \quad (1)$$

Varying the distribution of ω_i or altering the value of K moves the system between different dynamics. With sufficiently high K , all the oscillators will come together and oscillate at the mean frequency. At sufficiently low K the system is indistinguishable from a set of oscillators moving at their individual natural frequencies. Somewhere between these two extremes the system displays more interesting behavior [1] - see Figure 3 for an example.

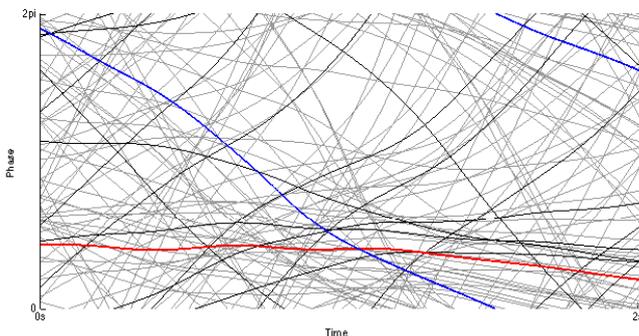


Figure 3: Kuramoto system of 66 weakly coupled oscillators. The different colors and shades of gray have no meaning - they are simply a visual aid to understanding the dynamics of the system. The natural frequencies are picked from a Gaussian distribution with mean zero.

To this simple system of weakly coupled oscillators can be added a variety of complications, the most natural of which is stochasticity, accounted for by adding a stochastic term to Equation 1. It is also possible to define a different coupling strength for each pair of oscillators, thus turning the constant K into a matrix K_{ij} . The third complication is to add a temporal offset into the phase difference calculation (sometimes this is approximated as a phase difference). Taking these three complications, the less trivial form of the Kuramoto model is:

$$\frac{d\theta_i}{dt} = \omega_i + \xi_i(t) + \sum_j K_{ij} \sin(\theta_j(t) - \theta_i(t - \tau_{ij})) \quad (2)$$

Several analytical results have been proven for the simpler cases, particularly for when there are an infinite number of oscillators, or where the coupling is just the constant K [30, 1]. In the latter case it is possible to reformulate the system of equations in terms of the circular mean, $\frac{1}{N} \sum_j e^{i\theta_j}$. This value is a complex number with an amplitude, r , and a phase ψ . Equation 1 now becomes:

$$\frac{d\theta_i}{dt} = \omega_i + Kr \sum_j \sin(\psi - \theta_j) \quad (3)$$

However, in the systems we are interested in the dynamics are likely to be closer to Equation 2, and will vary hugely depending on the characteristics of the two matrices and the properties of the stochasticity, not to mention the shape of the frequency distribution itself. We should also note that with only a single Fourier component (i.e. sine) the system's behavior is somewhat limited; adding a second component would significantly extend the model's behavioral range, even allowing it to approximate coupled Hodgkin-Huxley neurons [5].

2.4 Vicsek Particles

Of all the models discussed here, the Vicsek model has possibly the least direct relevance to the brain (it was introduced in the 1990s to describe flocking behavior in animals) however in this work we will be referring to it in other contexts.

The entities in the model are defined by their continuous-valued spatial coordinate and continuous-valued velocity direction (velocity magnitude is constant). At each timestep the direction of each particle is set to the mean direction of all the particles within a certain radius, R . A stochastic term is then added to this direction, before applying the velocity vector to update the particle's position [42]. Figure 4 illustrates the basic model.

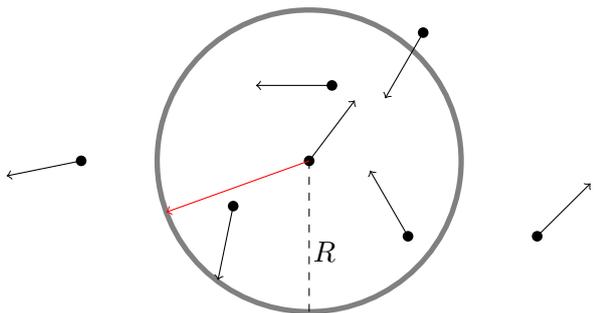


Figure 4: Seven Vicsek particles. The **red arrow** shows the mean direction of the four particles within a radius, R , of the central particle. Thus at the next timestep the central particle will move towards the south west (the exact angle will deviate from the red arrow according to the contribution of the stochastic term).

The parameters of the system are: the density of particles, ρ ; the magnitude of the stochastic term, η ; the magnitude of the velocity vector, v ; and of course the radius of influence, R . Note that the problem can be posed more neatly using three dimensionless parameters composed of these four quantities. In either case there will be regions of parameter space in which the interactions dominate, causing the particles to clump together, and other regions of parameter space in which the particles' motions are indistinguishable from independent random walks. As with the Kuramoto system there will be an intermediate region in which the behavior of the system is more interesting [42].

2.5 Integrate and Fire Neurons

Several models exist for describing the behavior of neurons. One of the simplest specifies membrane voltage dynamics using a differential equation that has to be integrated until the voltage reaches a certain threshold, at which point an action potential is said to occur and the variables in the equation are reset. The model takes its name from this integrate and fire regime. The differential equation itself is relatively simple, deriving from the law of capacitance [20]:

$$-C \frac{dV}{dt} = I_{ext}(t) + g(V(t) - V_L) \quad (4)$$

where C is the membrane's capacitance, $V(t)$ is the transmembrane voltage, and $I_{ext}(t)$ is the transmembrane current due to external sources (other neurons or artificially applied current). The term $g(V(t) - V_L)$ is actually an extension to the basic model, accounting for the current due to ions leaking

through the membrane under the presence of a non-equilibrium voltage.² Figure 5 demonstrates how the model functions.

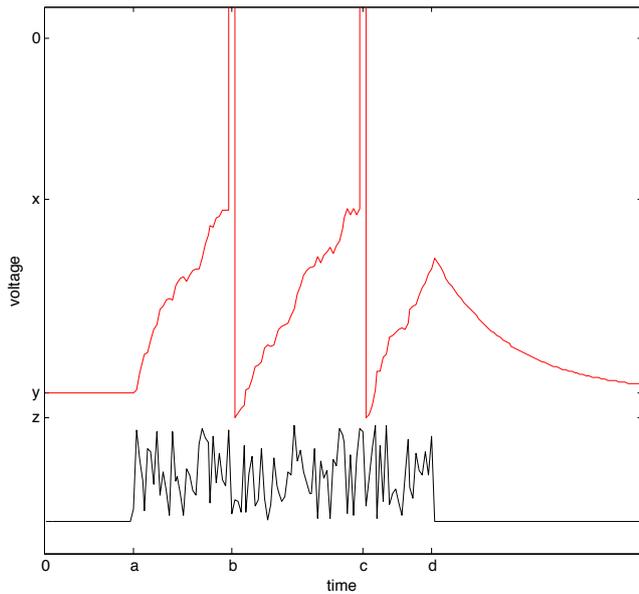


Figure 5: The **red line** shows the change of membrane voltage over time, simulated according to the integrate and fire model, where the **black line** gives the value of $I_{ext}(t)$, though not to scale. Between times a and b the differential equation is being integrated, then at time b an action potential is simulated because the voltage has passed the threshold x . Immediately after this event the voltage is reset to z , which is lower than its equilibrium value of y ($y = V_L$ in the equation). This process is repeated between times b and c . At time d the input is removed and the dynamics are purely driven by the leak term in the differential equation.

The single neuron so far described is not particularly interesting, the interest only arises when we simulate multiple such neurons and connect them up via the $I_{ext}(t)$ term. Designing this term can be quite complicated, taking into account connection strengths and the presence of different types of synapse, in some cases even modeling the concentration of key ions [20].

Although IF models have the potential to produce very complex dynamics when simulated in huge numbers, their relevance to understanding the brain's large scale dynamics is questionable as they do not abstract out any low level behavior. Despite this, we shall return to IF neurons briefly in the context of synchronization.

3 Mathematical Concepts

3.1 Networks

Turing's original formulation of a machine [37] was one in which computation occurred at a single location. However, as is the case with all the models described in the previous section, modern day computers and the brain perform computation at tens, hundreds, or even billions of locations simultaneously [40]. The Kuramoto and IF models in particular are similar to the brain in that their dynamics are governed by the strength of connections between the constituent entities. And, as has already been implied, the connections can be represented in a matrix, where the entry in the i th row and j th column gives the connection weight from entity i to entity j .

In order to understand the behavior of computational networks it is helpful to have some way of characterizing their properties so that different networks can be compared and analyses can be performed on particular classes of network. Many of the characteristics are defined for boolean networks, where the connections weights take the value 1 or 0 [34].

Small world networks Some networks are defined simply by the proximity of nodes in an N -dimensional space. The nodes may be arranged in a lattice - in a sense CAs and the Ising model are networks of this kind; or they may be scattered through space with a particular probability - the Vicsek model behaves like this over short time periods.

At the other extreme, the probability of two nodes being connected could be entirely random, with no influence from other connections in the network.

Between these two extremes there exist networks with mostly local connectivity, but some degree of random long range connections. Figure 6 shows these three classes of network on a 1-dimensional

²The equilibrium voltage is negative due to the competition between chemical gradients and electrical gradients setup by the various ion channels in the membrane. To be precise, it is the energy expended by a subset of these ion channels that maintains the voltage at its non-zero value.

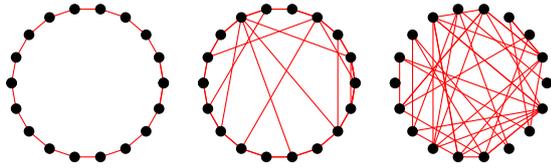


Figure 6: 1-dimensional periodic lattices of nodes. The **leftmost network** has only local connectivity, making it a true lattice graph; the **rightmost network** shows totally random connectivity; and the **central network** shows a small world graph. Adapted from [34].

periodic lattice.

Measures Figure 7 shows various simple network properties that can be calculated for individual elements or for pairs of elements. When describing an entire network it is generally necessary to calculate these properties for all possible individuals and pairs to produce a distribution that in turn can be used to summarize the system (e.g. the mean shortest path is known as the network’s diameter). In many cases though these simple measures will be insufficient to distinguish between two quite different types of network. One possibility [28] is to extend the concept of degree to apply not simply to the number of edges, but rather to the type and number of small subnetworks - for example there are 29 types of subnetwork/graphlets with five or fewer vertices. A second possibility is to look at the larger scale properties of the network: it may contain modules of vertices with high internal clustering; perhaps even a hierarchy of such modules.

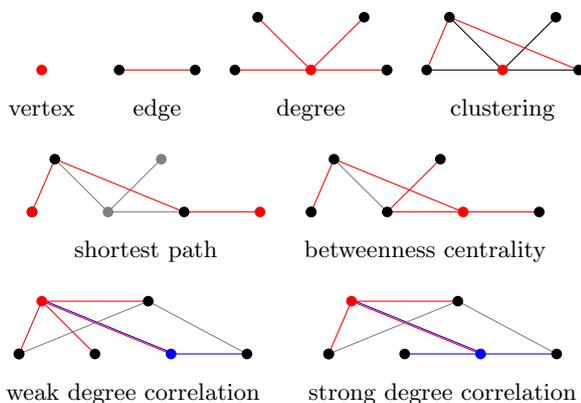


Figure 7: Formally, a network consists of **vertices** and **edges**. The **degree** of a vertex is defined as its number of edges, thus the node in red here has a degree of four. The **clustering coefficient** of a vertex is given by the fraction of its neighbors which are themselves connected, thus

for the red vertex shown here the clustering coefficient is $\frac{1}{3}$ because two out of a possible six neighbor-neighbor connections exist.

For every pair of vertices there exists a **shortest path** length defined by the smallest number of edges that must be traversed to get from one vertex to the other, thus the two red nodes here have a shortest path of three.

The **betweenness centrality** of a vertex is defined as the fraction of all the network’s shortest paths which go through it, thus the node in red here has a high centrality as almost all the shortest paths on the network pass through it.

For a pair of neighboring vertices, the **degree correlation** indicates whether the two vertices have a similar degree. In the network on the **left**, the red and blue vertices are neighbors but their degree correlation is relatively low because the red vertex has a degree of four and the blue vertex only has a degree of two. In the network on the **right** both red and blue vertices have a degree of three, and thus their correlation is much higher. Adapted from [34].

Scale free networks In many real world networks the distribution of degrees is said to follow a power law, meaning that the probability, $P(k)$, of any random node having degree k is given by $P(k) = k^{-\gamma}$, with the exponent γ being roughly equal to 2 [34]. Such networks can be generated by starting with a random graph and then, in an iterative manner, preferentially adding edges to nodes with high degree.

3.2 Criticality

Several times during our discussion of complex systems we noted that there may exist a particular region of parameter space which produces significantly more interesting behavior than all other regions - as depicted in Figure 8. Here we explain in a little more detail what this means, before in later sections considering its implications for the brain.

The Ising model is often used to introduce criticality [12, 4] and we shall continue that tradition here. At a critical temperature in the Ising model the system is neither totally homogeneous nor is it simply pure noise. Instead its behavior is rather interesting in the sense that there are dynamic correlations across space (dynamic correlation is equivalent here to covariance, i.e. for a pair of elements it is maximal if they simultaneously vary from their means with the same sign). One particular temperature is said to be ‘critical’ because it maximizes these correlations at every spatial scale.

This brings us to an important recurring feature of systems operating at a critical point: they tend to exhibit many ‘power laws’ of the form $y = x^{-r}$, where y is some quantity measured at x with a

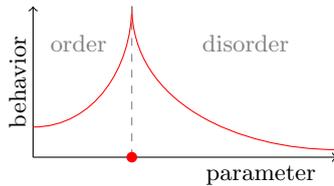


Figure 8: Cartoon depicting criticality. In the Ising model the parameter is temperature and the behavior is the dynamic correlation. The idea is that at some critical value of the parameter the behavior suddenly becomes very special. Adapted from [4].

constant exponent r . For the Ising model this relationship holds if x represents the distance between a pair of elements and y is the dynamic correlation [4].

An equation of this form was presented in the previous section in the context of scale free networks. This is no coincidence - power laws are inherently scale free because multiplying x by some constant, s , does not alter the exponent: $y = (sx)^{-r} = s^{-r}x^{-r}$ [25]. The implication of this is that the behavior of y looks similar at all scales of x .

In all the models introduced here a parameter must be set to a particular value for the system to be critical and demonstrate power laws. However in some classes of system the same power laws exist at all regions in parameter space (or rather there are no important parameters in the system). Such systems are said to exhibit self-organised criticality [2].

3.3 Information Theoretic Measures

Unlike all the models described thus far, the brain is not homogenous - different regions have different local connectivity and their constituent neurons behave in different ways. This variation means that for any given computational task some regions will be more successful and efficient than others. Indeed each class of problem will demand a different balance between optimizing speed, energy use, neural ‘real estate’, and accuracy [21, p. 470][17, 40, 22].³ Characterizing the behavior of a computing model or of the brain itself therefore requires a range of measures that capture different aspects of the information processing [7].

Entropy and conditional entropy ⁴For a data-producing entity, i.e. an experiment or more generally a random variable, the quantity of information provided by each new piece of data depends on how probable the new data was: observing an unsurprising result provides little information relative to observing a surprising result. Mathematically, the information obtained when X is observed to have the value x is defined as $-\log p(X = x)$ or just $-\log p(x)$, where the convention is for the logarithm to be taken in base 2.

The expected information provided by X is an important quantity known as the Shannon entropy. It is given by:

$$H_X = \mathbb{E}_{x \in X} [-\log p(x)] = -\sum_x p(x) \log p(x) \quad (5)$$

For two data-producing entities, X and Y , the joint entropy is given by the joint probability distribution:

$$H_{X,Y} = -\sum_{x,y} p(x,y) \log p(x,y) \quad (6)$$

If Y is observed to have the value y , but X has not yet been observed, we can use the joint probability distribution, $p(x,y)$, to calculate $p(x|y)$, and thus calculate a new entropy for X :

$$H_{X|Y=y} = -\sum_x p(x|y) \log p(x|y) \quad (7)$$

Taking the expectation over all Y gives the conditional entropy:

$$H_{X|Y} = \mathbb{E}_{y \in Y} \left[-\sum_x p(x|y) \log p(x|y) \right] = -\sum_y p(y) \sum_x p(x|y) \log p(x|y) = -\sum_{x,y} p(x,y) \log p(x|y) \quad (8)$$

³In recent years the architecture of consumer electronic devices has become more like the brain in this regard: the central processing unit can delegate arithmetic intensive tasks to a special floating point unit, and it can delegate graphics-related tasks to a separate processor, the constituent parts of which are optimized for performing particular stages in the graphics pipeline.

⁴A concise introduction to information theory is given in [24], much of which is paraphrased here.

Mutual information If Y does little to inform us about X , we will find that $H_{X|Y}$ is almost as large as H_X . Conversely if Y greatly improves our prediction of X we will find that $H_{X|Y}$ is much smaller than H_X . The difference is known as the mutual information of X and Y :

$$H_{X;Y} = H_X - H_{X|Y} \quad (9)$$

Entropy rate and active information storage If the individual pieces of data generated over time by X , that is x_1, x_2, x_3, \dots , are independent, then x_{i+1} will have zero mutual information with $\{x_i, x_{i-1}, x_{i-2}, \dots, x_{i-k+1}\}$, which we abbreviate to $x_i^{(k)}$. If however there is some dependence across time, then the mutual information will be non-zero and (so long as the process is stationary) we can define a quantity - ‘the entropy rate’ - for the limit of the conditional entropy:

$$H_{\mu X} = \lim_{k \rightarrow \infty} H_{x_{i+1}|x_i^{(k)}} \quad (10)$$

Using the fact⁵ that $H_{A|B} = H_{A,B} - H_B$ we can express this as the limit of $H_{\mu X}(k)$ where:

$$H_{\mu X}(k) = H_{x_i^{(k+1)}} - H_{x_i^{(k)}} \quad (11)$$

Related to the entropy rate via Equation 9 is the ‘active information storage’:

$$A_X = \lim_{k \rightarrow \infty} A_X(k) \quad \text{where } A_X(k) = H_{x_{i+1};x_i^{(k)}} \quad (12)$$

Figure 10 shows how active information storage can be evaluated locally, rather than averaged. And Figure 24 explains how using information from the past can alter the information content of future data, potentially making the active storage negative.

Transfer entropy The entropy rate tells us to what extent X is dependent on its past, and mutual information tells us the extent to which X and Y are dependent on one another. Combining these two ideas, we can calculate the rate at which information present in X is being transferred to Y [7]:

$$T_{X \rightarrow Y} = H_{y_{i+1}|x_i^{(k)}, y_i^{(k')}} - H_{y_{i+1}|y_i^{(k')}} = \sum_{x^{(k'+1)}, y^{(k+1)}} p(y_{i+1}, y_i^{(k)}, x_i^{(k')}) \log \frac{p(y_{i+1}|y_i^{(k)}, x_i^{(k')})}{p(y_{i+1}|y_i^{(k)})} \quad (13)$$

This accounts for y 's dependence on past y values, but it cannot test whether X and Y are really communicating - a third party could be transferring information to X first and then Y .

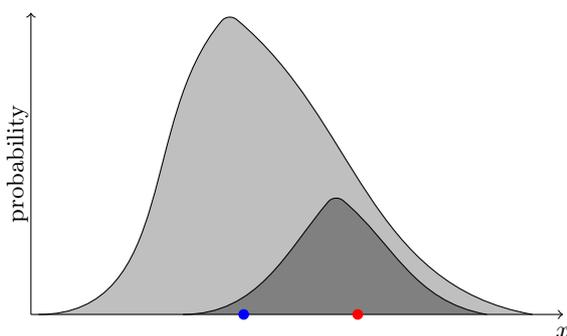


Figure 9: Cartoon depicting active information storage. In **pale gray** is the naive probability distribution for x_{i+1} . Observing the **red** value is not particularly likely, but observing the **blue** value is reasonably likely. Suppose we know the values of $x_i^{(k)}$, and can update the probability distribution to take this into account. We might now have something like that shown in **dark gray** - strictly speaking both distributions should integrate to 1, however here we depict the conditional distribution as being smaller to reinforce its meaning. Note how the new information makes it less likely that we observe **blue** and more likely that we will observe **red**.

Excess entropy Although the entropy rate tells us the extent to which x_i is dependent on all past values of X , it does not capture how x_{i+1}, x_{i+2}, \dots will also depend on the previous values of X . This can be calculated using the following infinite series:

$$E_X = \sum_{k=0}^{\infty} [H_{\mu X}(k) - H_{\mu X}] \quad (14)$$

⁵The proof of the identity is arrived at from Equation 8 by expressing $p(x|y)$ as $p(x)/p(x,y)$ and expanding out the logarithm.

where the term $[\cdot]$ essentially gives the excess uncertainty of a naive observer who only knows the k last values of X . E_X can also be expressed as the mutual information between the past and the future [24]:

$$E_X = \lim_{k \rightarrow \infty} E_X(k) \quad (15)$$

where $E_X(k)$ is the mutual information between the k past and k future values of X . As shown in Figure 10, E_X quantifies how much information is being stored for use at some point in the future. Note however that the information is not stored in individual cells in the CA, rather it is distributed over small structures some of which are static and others of which travel along the axis of the CA.

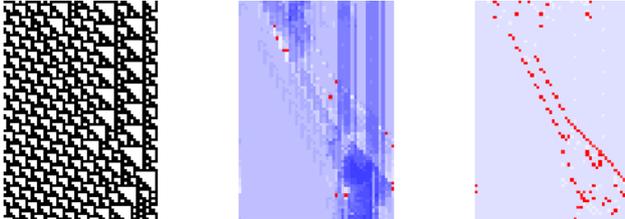


Figure 10: The **left panel** shows a 1-dimensional CA using rule 110 to iterate over several generations, shown going down the page. The **central panel** shows the excess entropy evaluated for the corresponding cell-time pairs. Blue values are positive, red values are negative. The **right panel** shows the active information storage for the same cell-time pairs.

4 Biology of the Brain

The human brain is evolution's current solution to a somewhat poorly defined problem. As with all other examples of evolution, the task is essentially to optimize fitness, but for the brain the optimality will depend significantly upon the environment. Crucially, the fitness of the brain is not judged on its average daily ability, but also on its facility for dealing appropriately with rare, sudden, or novel events [33]. Learning from past information is fundamental to the brain's success, but it is not enough to simply store this information, it must be assimilated in a fashion that permits retrieval in a timely manner, in many cases requiring the information to be accessible almost instantaneously [36].

In this section we briefly explore how evolution has addressed these design requirements, focusing in particular on the results at the mesoscopic scale.

4.1 Mesoscopic Architecture of the Brain

The main tools for examining the functional architecture of the brain at the mesoscopic scale are MRI, MEG and EEG.

At its simplest, MRI can be used to measure the electromagnetic environment of a particular chemical species in a sample. Modern MRI machines are capable of millimeter/second resolution, and in the case of fMRI can be used to measure the oxygenation level of blood, a proxy for neural activity [10]. A second use of MRI is for DSI, which measures anisotropies in the movement of water molecules. These anisotropies reveal the orientation of axonal bundles in each voxel of white matter, allowing a computer to trace axonal tracts from one region of gray matter to another.⁶ Figure 11 shows the human connectome at the mesoscopic scale which was obtained using DSI, it also shows the correlation between regions observed using fMRI.

There is some debate as to what network properties the connectome exhibits. It is often claimed that it has the small world property, is scale free, and that wiring lengths are minimal. However there is some doubt as to whether any of these claims apply at the mesoscopic scale, or indeed at all [15, 17]: validating power laws is difficult due to the requirement of measuring vertex degree over several magnitudes; validating the small world property is difficult due to the requirement of tracing long range connections that may be very weak; and validating the optimal wiring claim is difficult due to not knowing the full constraints placed on the network.

An important empirical observation from fMRI studies is that regions which are strongly connected tend to have correlated activity: this can be seen from the similar structure of the two matrices in Figure 11. In addition, there is plenty of evidence - both from fMRI and recently from MEG [6, 16] - that the mesoscopic brain network consists of about eight partially independent subnetworks, identified

⁶This is described in detail in [15], but we provide a summary here. Each white matter voxel is seeded with several fibers that iteratively extend out to neighboring voxels, changing their direction only gradually. Fibers that reach dead ends (due to requiring sudden changes in direction) are discarded. Note that DSI produces a spectrum of directions in each voxel, meaning that in a single voxel there can be axons orientated in multiple directions.

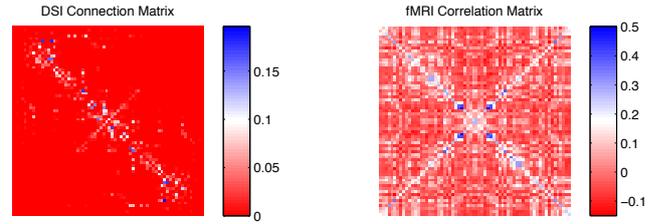
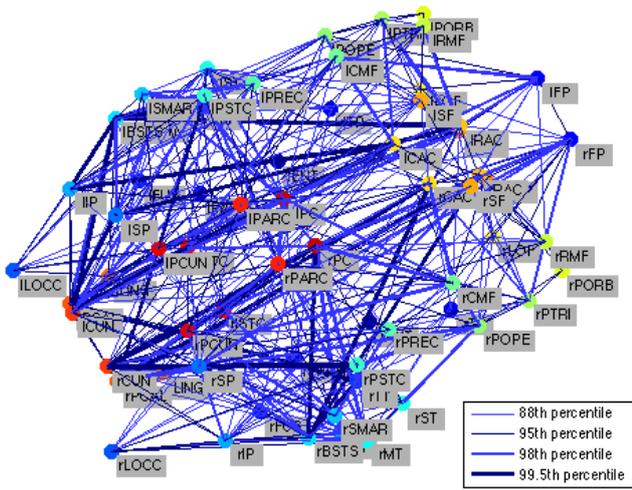


Figure 11: **Left.** The human brain at the mesoscopic scale: a network of 66 nodes. **Legend** shows how the line color and weight correspond to the connection strength, where percentiles are calculated over the whole connection matrix. Key to labels is not provided here, see original source for details. **Above.** Connection matrix used to produce the network plot, and functional connectivity matrix as measured by fMRI correlations. Data from [15, 16] and adapted from [10].

by their particularly strong internal correlations. These subnetworks seem to correspond to the major tasks that the brain deals with, such as vision, sensorimotor activity, and abstract thought. Figure 12 shows evidence for the most well known of these subnetworks.

Another avenue of research is to examine neurological pathologies in which the brain’s mesoscopic connectome varies significantly from the norm: schizophrenia, dementia, attention deficit disorder, and epilepsy have all been examined in this context as they may be examples of such a pathology [34].

4.2 Brain Waves

One of the most fundamental features of the brain’s activity is that from a very early age until death it remains reasonably constant: there are no periods of ‘silence’, neither are there periods of intense ‘storms’ (except in the case of epilepsy). Achieving the right balance is important for the various reasons alluded to in the introduction to this section.

The most important mechanism by which this balance is attained is by the extensive use of temporal oscillations which allow the system to safely cycle between alternating periods of low and high activity without becoming stuck in either state. These temporal oscillations exist over a wide range of frequencies, and in the > 1Hz range have the following names [6]:

range	1-4Hz	4-8Hz	8-13Hz	13-30Hz	30-50Hz
name	delta	theta	alpha	beta	gamma

Using EEG, MEG, or invasive electrophysiological recording it is possible to record these waves [8] and then perform Fourier, wavelet or Hilbert analysis to label each time point with a phase and a magnitude in each of the frequency bands. Examining the spatial correlations in these values - or indeed in the time domain signal - can provide information about synchronization of brain regions. Examining the frequency-power distribution reveals something close to a power law with exponent equal to 1 [9]. We shall explore the implications of this power law and the existence of synchronization in the next

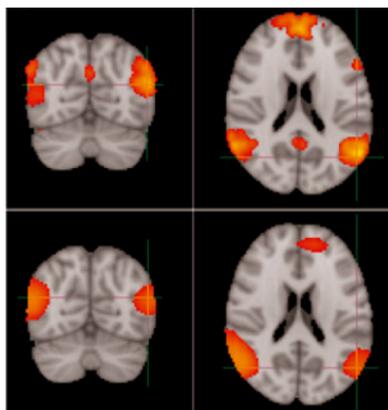


Figure 12: The **top row** shows data obtained from fMRI, and the **bottom row** shows data obtained from MEG. In both cases the highlighted regions display highly correlated activity (they were identified using independent component analysis). This particular brain subnetwork is known as the default mode network because it is most active when the subject is at rest, not performing any obvious mental task. Taken from [6].

section, in the remainder of this section we give a brief overview of two of the known mechanisms by which oscillations are generated and then discuss a few other related empirical findings.

As described mathematically in Figure 13, the membrane voltage of individual neurons tends to oscillate naturally due to the interactions of ion channels, an effect that can be measured by patch-clamping a cell in vitro [8]. This natural frequency depends on various properties of the cell and as such appears to range from the theta band to the gamma band depending on cell type.

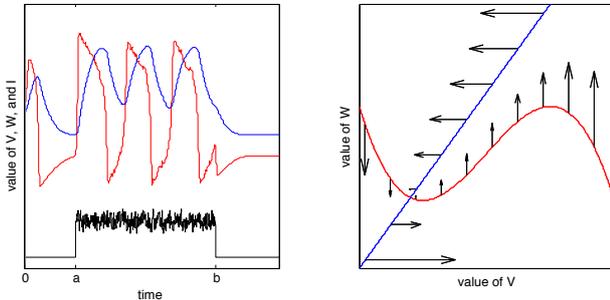


Figure 13: The famous Hodgkin-Huxley neuron is more complex than the IF neuron described earlier but it has the advantage of producing realistic dynamics without any awkward discontinuities. Helpfully, the full Hodgkin-Huxley model can be approximated as a system of just two variables - the FitzHugh-Nagumo neuron [21, pp. 173-184]. This

model has three parameters: a , b , and ϕ , as well as the time series for the externally applied current, I :

$$\frac{dV}{dt} = V - \frac{V^3}{3} - W + I \quad \frac{dW}{dt} = \phi(V + a - bW) \quad (16)$$

The **left panel** shows how V , in **red**, and W , in **blue**, evolve over time. The input, I , is shown in **black** and is non-zero between times a and b .

The **right panel** shows the phase space of V and W for fixed I . Along the **red line** $\frac{dV}{dt} = 0$, and along the **blue line** $\frac{dW}{dt} = 0$. The **black arrows** show the non-zero component of the derivative evaluated along the two lines. The key observation is that there is one stable point (the intersection of the two lines), and that in the presence of noise or varying input the neuron will cycle around this point alternating between high voltage (a spike) and low voltage. Plots adapted from [21, pp. 173-184].

A second mechanism for generating oscillations is the inhibition-excitation loop [3] which has been extensively examined in the hippocampus, but does exist elsewhere too. Figure 14 shows two possible forms of this loop, both of which are supported by evidence: global reduction of inhibition (using the GABA_A antagonist bicuculline) stops gamma oscillations forming when they otherwise would under electrical stimulation or pharmacological stimulation (using carbachol or kainate); global reduction of excitation (using the AMPA antagonist NBQX) also stops gamma oscillations forming under pharmacological stimulation (again using carbachol); however global reduction of excitation using a different pharmacological method (the AMPA antagonist GYKI 53655) does not have an impact on the gamma oscillations formed under pharmacological stimulation (using kainate) [3].

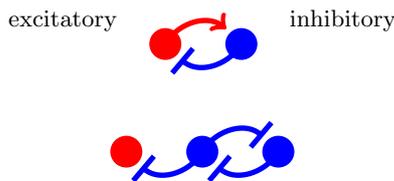


Figure 14: Two possible inhibition-excitation loops that would generate gamma oscillations. Each circle represents a neuron or pool of neurons with recurrent connectivity. See main text for explanation of supporting evidence. Adapted from [3].

At this point we note that EEG, MEG, and even invasive extracellular LFP recording do not measure spiking rates directly. The LFP is thought to reflect the total level of synaptic activity [8] occurring within approximately a $250\mu\text{m}$ radius [19]. As depicted in Figure 15, the strength of the signal depends on the dipole-like nature of neurons and the extent of their misalignment. EEG at the scalp, is thought to be capable of recording the quasi-simultaneous spiking of anything over 10^8 neurons, an event known as an avalanche [25]. In a variety of experiments the duration and magnitude of these avalanches have been shown to approximate a power law, though there is some disagreement as to the value of the exponent [25].

In addition to measuring spatial coherence of oscillations in EEG or MEG signals it is also possible to examine heterogeneities in power at the mesoscopic scale, and heterogeneities with respect to developmental stage and state of consciousness [38, 26].

Finally, using MEG, EEG, or fMRI, it is possible to examine the series of events that occur in the brain as part of the response to a particular stimulus, presented to the subject multiple times [38].

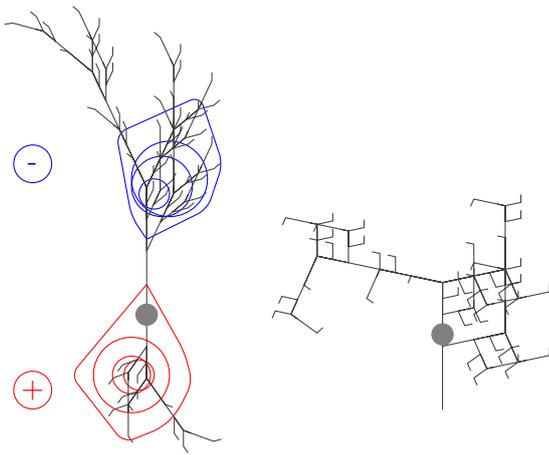


Figure 15: On the **far left** is a cartoon of a neuron in **gray**. Shown in **blue** is the effect of the depolarization current in the dendrites, which causes a buildup of negative charge in the extracellular space (due to calcium and sodium ions flowing into the neuron). And shown in **red** is the effect of the balancing current that results from this dendritic depolarization. This process sets up a dipole.

On the **right** is a neuron where the axons and dendrites are not so well separated. Here positive and negative charges will tend to cancel locally and may not produce a significant dipole. Of course if only one section of the dendrites is active the cell will be able to act more like a true dipole. Even if the processes of a single neuron are separated, as in the example on the left, neighboring neurons will tend to cancel out the charges if they are not vertically aligned. This misalignment becomes more pronounced in the larger mammals and hence results in a weaker extracellular electric field. Adapted from [8].

4.3 The Role of Synchronization

As already mentioned, it is assumed that the brain employs a range of mechanisms for performing its computations. For example [21, pp. 471-472], it is thought that the multiplication of two numbers, x and y , may be performed by either (a) appropriate positioning of excitatory and inhibitory synapses on a single dendrite; or by (b) a simple summing of two inputs corresponding to $\log x$ and $\log y$, followed by an exponentiation of the result; or by (c) assuming x and y are encoded by the rates of two independent Poisson processes and using the fact that $p(x \wedge y) = p(x)p(y)$; or by (d) taking the difference of the mean rate from two populations of cells, one of which calculates $(x + y)^2$ and the other of which calculates $(x - y)^2$, where the squaring is achieved if the population consists of cells that have a linear response to inputs above a threshold picked from a uniform distribution.

Although there is some evidence for the above mechanisms, they do not deal with the implications of neurons being oscillatory; more generally they fail to approach the question of time-varying thresholds which may be entirely stochastic or may show a strong dependence on the history of synaptic activity on the scale of 250ms [14]. The ‘Communication Through Coherence’ hypothesis states that these effects are important and give rise to short windows of time in which a neuron will be maximally sensitive to input. Further, if one pool of neurons needs to communicate with a second pool it will deliver maximum effect by not only targeting this optimum excitability window, but also by firing in a synchronized burst. Given, as we have already said, neurons tend to oscillate at about 40Hz, it seems plausible that this maximum information transfer can be achieved if the two pools of neurons lock their phases, i.e. their oscillations are coherent. This idea is summarized in Figure 16

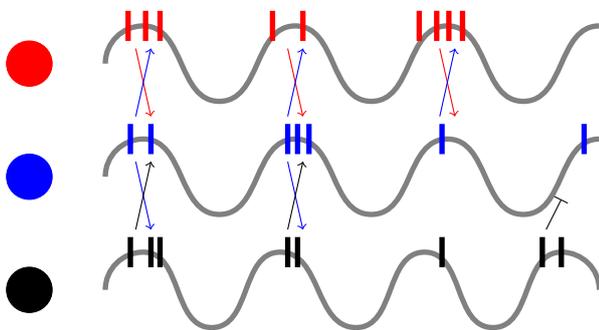


Figure 16: In **red**, **blue**, and **black** are three pools of neurons. The **gray sinusoidal curves** show the subthreshold voltage oscillations in each pool, and the **vertical bars** represent spikes from the corresponding pool. Initially all three pools are in phase, and thus their bursting activity coincides and has maximal impact on the target pool. After two cycles however, the black pool is no longer in phase (i.e. it was never coherent) and it may actually begin to inhibit activity in the blue pool. Adapted from [14].

Communication through coherence is potentially a useful mechanism for generating temporary high-bandwidth connections between distant regions. Note however that the hypothesis does not explain how regions fall in and out of coherence with one another, only that during periods of coherence they will be better placed to transfer information. Demonstrating that two regions are phase locked in the gamma band is not too difficult as gamma is easily observable, however showing that information transfer is correlated with the extent of this coherence is rather taxing. One attempt [45] used electrophysiological recordings from cat and monkey to show that correlations of gamma band power

are related specifically to coherence in the gamma band, and that they may be causally related since the statistical dependence is strongest about 5ms after synchronization. However the analysis did not examine the transfer of any information beyond the single value of power so it is difficult to claim that this truly reflects the extent of network communication. In response to this experimental result a computational model [7] was produced that consisted of a network of IF neurons, the advantage of the model being that a finer level of analysis could be performed - the authors measure the total population activity and conclude that transfer entropy is greatest during coherence and is highest when the gamma band power is strongest.

Counter to the communication through coherence hypothesis, there is in vitro evidence [23] that significant communication between populations can occur in the absence of total population firing rate oscillations (and even in the absence of information content in the source population's total firing rate). This experimental observation is matched by a recent [29] piece of modeling work that takes a network of IF neurons and examines its ability to distinguish between different externally supplied stimuli.

The existence of distinct windows of low and high excitability may serve a different computational role [22]: using a lattice of IF neurons it has been shown that low excitability is better for distinguishing between specific temporal spike patterns, whereas high excitability is better at distinguishing between different spiking rates. A related piece of analysis performed on this model was to examine the difference in network activity produced by pseudo-infinitesimal changes to the input. These differences are quantified by plotting against time the mean Euclidean distance between the network activity for the different inputs, and then fitting an exponential curve. At a critical value of the network-architecture parameters, it was found that the (Lyapanov) exponent of the curve was close to one. Perhaps unsurprisingly this point in parameter space also produced the best networks for classifying input timeseries.

A simpler line of enquiry when considering how the population responds to a range of inputs, is to characterize the input and output purely by their total magnitude in the population. A reasonably complex CA has been used to demonstrate that at a critical value of parameters, variation in input magnitude will be maximally mapped to variation in output [20]. In response to this computational finding, rat brain slices were stimulated with a range of currents and the magnitude of the population response was recorded [32]. The resulting input-to-output map did indeed cover a wide range of inputs, and was tentatively considered to be the critically tuned curve given that pharmacological alteration of the slices reduced its dynamic range.

5 New Work

5.1 Discussion of the model

In this work we will be modeling the brain at the mesoscopic scale using a Kuramoto system of weakly-coupled oscillators. This has previously [10] been done using Equation 2 and the connection matrix, K_{ij} , shown in Figure 11. Noise and oscillator frequencies were picked from Gaussian distributions, where the standard deviations, σ_n and σ_f , were treated as parameters. The delay matrix, τ_{ij} , was derived from the DSI data set by averaging the fibre lengths between regions and assuming a constant transmission speed. The transmission speed itself was not specified, rather it was parameterized implicitly by treating the mean delay, $\langle\tau\rangle$, as a parameter. Similarly, the connection matrix was multiplied by a fourth parameter, k .

Simulations began by setting oscillator phases to random values picked from a uniform distribution on the interval $[0, 2\pi]$. The model was then run for a few timesteps without any coupling, before switching to the fully coupled mode (this was necessary due to the delay matrix requiring each oscillator to have a history). In practice these details were not important because the analysis was not applied to the first 20 seconds of data, the only section notably effected by the initial conditions.

The simulations were run using the simplest type of numerical integration - the Euler method. The error introduced by this method was minimized by using a small timestep, but deviation from the exact solution was in reality not a problem as the system was intended to be stochastic and chaotic. The timeseries of oscillator phases produced by the simulation was downsampled to 1ms, and then converted to a simulated firing rate by taking the sine of the phase. Finally, the Pearson correlation

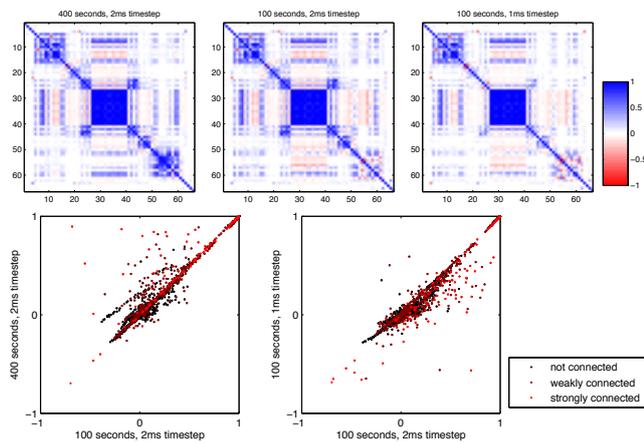


Figure 17: The **top** row shows the correlation matrices for the same simulation run three times using $(k=18, \langle \tau \rangle = 11, \sigma_f = 3, \sigma_n = 2)$ but with two different timesteps and for two different durations. In the **plots below**, each point corresponds to the correlation between a pair of oscillators, ie. a particular element from two of the above matrices. The 5% most strongly connected oscillator pairs are in **bright red**, the remaining connected pairs are in **dark red** and the non-connected pairs are in **black**. Even though the simulation is stochastic the difference between the matrices is clearly small, especially for the strongly connected pairs of oscillators.

between all pairs of oscillators was computed, producing a correlation matrix.

In [10] the timestep used was 0.1ms and simulations were run for a duration of 400 or 1200 seconds; here however we use a timestep of 0.2ms and only run simulations for 100 seconds (as in [10], the first 20 seconds are discarded). The difference in the correlation matrices is marginal as can be seen in Figure 17. The simulations for this work were run in MATLAB using specially written code, note however that the author was also given access to the original code used in [10].

The claim in [10] is that by using oscillators with a mean frequency of 60Hz the simulated system represents the brain's gamma-band driven behavior. However, apart from $\langle \tau \rangle$, the analysis does not involve any other indicators of temporal scale and should thus be interpreted with caution. The only counter evidence to this being that the chosen value of $\langle \tau \rangle$ does correspond to an appropriately scaled transmission velocity, however the results are not entirely different when much larger values are used.

The second major claim in [10] is that there is a particular pair of k and $\langle \tau \rangle$ for which the system has similar behavior to the fMRI signal, as defined by the matrix in Figure 11. To support this claim the simulated matrix is compared to the fMRI matrix using a Pearson correlation and using a sum of squared errors. This is done using only the values for connected oscillator pairs. Reducing the data in this way has two effects: it increases the contribution of the connected pairs, and it completely ignores any correlations in the non-connected pairs. While the former effect is potentially desirable, the latter is not: the system should be sensitive to interactions over more than a single network edge, with the interactions potentially being evident in the correlation matrix - see Figure 18.

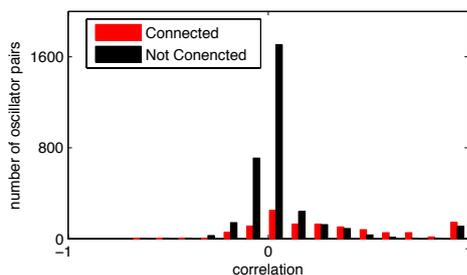


Figure 18: This plot shows the two potential effects of discarding the correlations for non-connected oscillator pairs: the proportion of non-zero values will increase, but a significant proportion of non-zero values will be discarded in the process. The parameter values used here were $(k=18, \langle \tau \rangle = 11, \sigma_f = 3, \sigma_n = 2)$.

Ignoring the details of the comparison between the fMRI matrix and the Kuramoto matrix, a major problem still stands: at the optimal parameter values chosen in [10] the Kuramoto matrix correlates strongly with the connection matrix used to produce the simulation. Indeed it is accepted that the connection matrix and fMRI matrix are themselves highly correlated. The implication of these observations is that the Kuramoto simulation may not be adding much to our understanding of the mesoscopic brain. See Figure 19.

5.2 Comparison with MEG data

So far the Kuramoto model has only been compared to fMRI data. Although this comparison was not entirely satisfactory we shall continue to examine the model here in a similar fashion. MEG data, unlike fMRI data, provides a high temporal resolution and is directly indicative of electrical activity

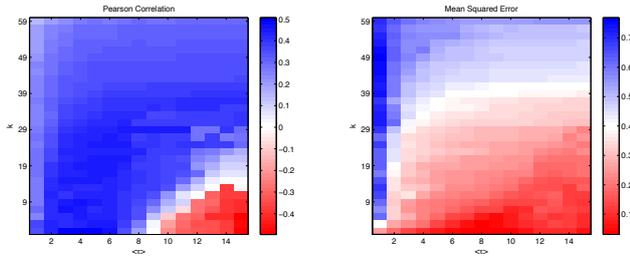


Figure 19: Comparison of the Kuramoto correlation matrix with the connectivity matrix at a range of $(k, \langle\tau\rangle)$ values. Note that at $(k=18, \langle\tau\rangle=11)$ the correlation is high and the error is low. In [10], there are several levels of data processing applied to the raw Kuramoto data to make it more fMRI like, however the resulting correlation and mean squared error values at $(k=18, \langle\tau\rangle=11)$ are very similar to those shown here.

rather than being a proxy via metabolic and hemodynamic responses [6]. Thus it may reveal different aspects of the mesoscopic behavior of the brain.

The main disadvantage of MEG is that it only samples the magnetic field outside the skull, and thus, much like a shadow falling on a wall, it can only ever provide an estimate of the true volumetric properties from which it is projected [6]. In some usage cases the exact location of a source is not particularly important and analysis can be performed in sensor space, however in this work we wish to know the level of activity at 66 specific points inside the brain so we must attempt some form of volumetric reconstruction.

Figure 20 shows the geometric setup of the system. We now explain, in brief, a procedure called Synthetic Aperture Mapping [41] which is used to project from sensor space into source space. Rather than jumping to a general formalism here, we begin by considering a system with only 3 sensors and show the required calculation. The projection process is carried out independently for each point of interest, so here we need only concern ourselves with a single point, θ .

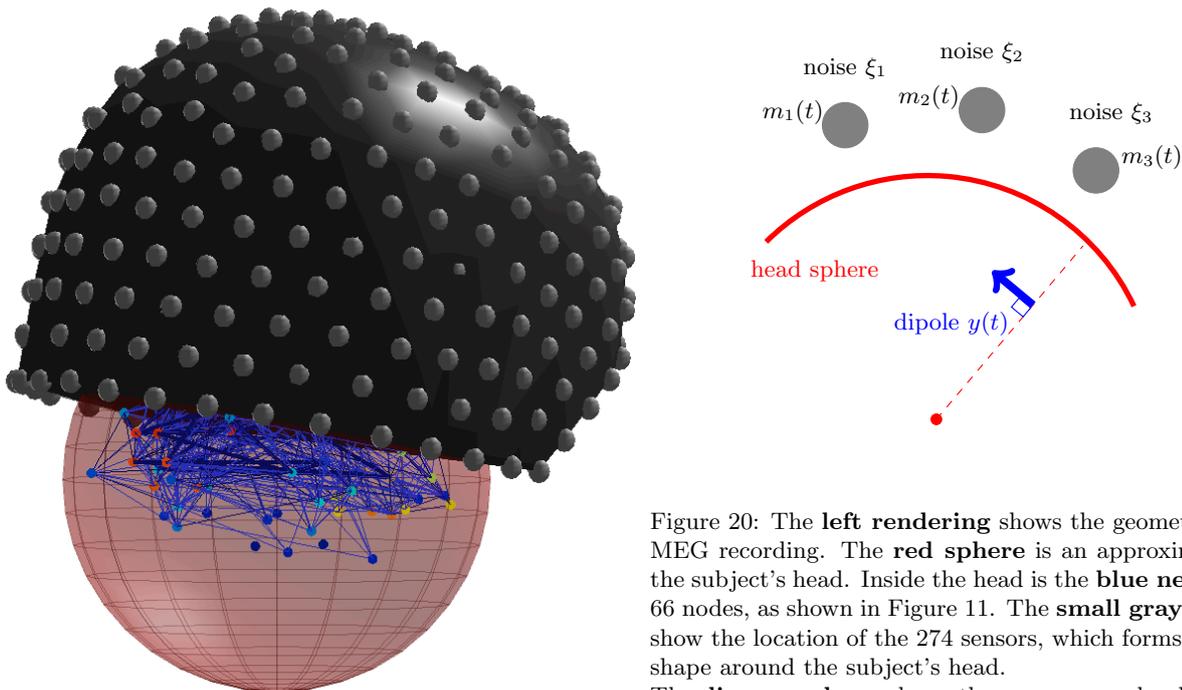


Figure 20: The **left rendering** shows the geometry of the MEG recording. The **red sphere** is an approximation of the subject's head. Inside the head is the **blue network** of 66 nodes, as shown in Figure 11. The **small gray spheres** show the location of the 274 sensors, which forms a helmet shape around the subject's head.

The **diagram above** shows three sensors and a dipole, see the main text for a full explanation.

The task is to find an optimum vector of weights, $W = [w_1, w_2, w_3]$, which will be used to generate the timeseries, $y(t)$, at the point θ using the sensor timeseries, $M(t) = [m_1(t), m_2(t), m_3(t)]$:

$$y(t) = w_1 m_1(t) + w_2 m_2(t) + w_3 m_3(t) \quad (17)$$

$y(t)$ corresponds to the magnitude of a dipole at θ , the orientation of which we must specify or find by minimising some quantity, to be explained later. The optimality of the weights is defined here as meaning that $y(t)$ should have minimal power, i.e. we wish to minimize $\mathbb{E}[y(t)^2]$. Note that this minimization is only valid if the subject maintains a relatively constant mental state throughout the specified period.

$$P = \mathbb{E}[y(t)^2]$$

$$\begin{aligned}
 &= \mathbb{E} [(w_1 m_1(t) + w_2 m_2(t) + w_3 m_3(t))^2] \\
 &= \mathbb{E}[w_1 w_1 m_1(t) m_1(t)] + \mathbb{E}[w_1 w_2 m_1(t) m_2(t)] + \dots + \mathbb{E}[w_3 w_3 m_3(t) m_3(t)] \\
 &= w_1 w_1 \mathbb{E}[m_1(t) m_1(t)] + w_1 w_2 \mathbb{E}[m_1(t) m_2(t)] + \dots + w_3 w_3 \mathbb{E}[m_3(t) m_3(t)] \\
 &= w_1 w_1 C_{11} + w_1 w_2 C_{12} + \dots + w_3 w_3 C_{33}
 \end{aligned} \tag{18}$$

where C_{ij} is the covariance matrix for the sensor timeseries (we require that all sensors have mean zero). At the minimum value of P all partial derivatives are zero, thus we have a system of three linear equations:

$$\begin{aligned}
 0 &= \frac{\partial P}{\partial w_1} = 2w_1 C_{11} + 2w_2 C_{12} + 2w_3 C_{13} \\
 0 &= \frac{\partial P}{\partial w_2} = 2w_2 C_{22} + 2w_1 C_{12} + 2w_3 C_{23} \\
 0 &= \frac{\partial P}{\partial w_3} = 2w_3 C_{33} + 2w_1 C_{13} + 2w_2 C_{23}
 \end{aligned} \tag{19}$$

We do not want the null solution for W , rather we want the closest answer that meets the following constraint:

$$1 = w_1 b_1 + w_2 b_2 + w_3 b_3 \tag{20}$$

where $B = [b_1, b_2, b_3]$ specifies the sensitivity of the three sensors to the dipole. The calculation of these sensitivities uses standard Maxwellian physics, taking as input the orientations of the dipole and sensors, as well as the distance separating them. The volume is considered to be a vacuum with no currents. It is possible to specify the orientation of the dipole, but it can also be found by maximising the signal to noise ratio over all orientations. The power of the signal is given in Equation 18, and the power of the noise is simply:

$$N = w_1^2 \xi_1^2 + w_2^2 \xi_2^2 + w_3^2 \xi_3^2 \tag{21}$$

where $[\xi_1^2, \xi_2^2, \xi_3^2]$ specifies the magnitude of the noise at each sensor.

The answer to the constrained problem given by Equations 19 and 20 is obtained by using a Lagrange multiplier and then completing the square [39] (this is not trivial):

$$W = \frac{1}{B^T C^{-1} B} B^T C^{-1} \tag{22}$$

In this work we perform the minimization using the freely available MATLAB toolbox Field-Trip [27]. Dipole orientations are found by maximizing the signal to noise power as described above, with the additional constraint that the dipole must lie perpendicular to the radius of the sphere shown in Figure 20. Figure 21 shows the results for 60 seconds of data from a single subject at rest with eyes closed.

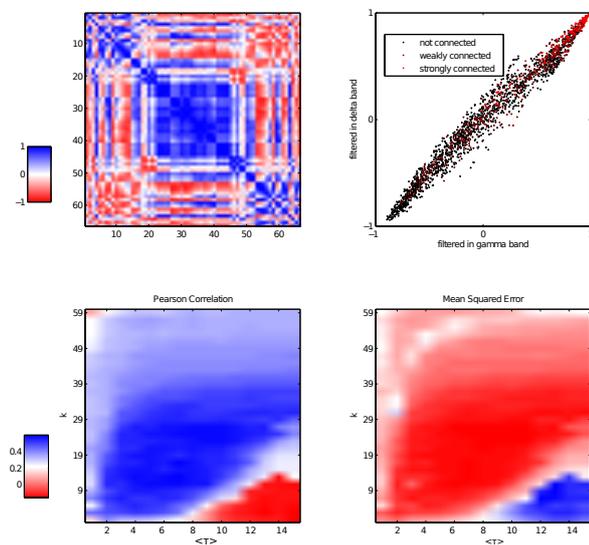


Figure 21: The **top left** plot shows the correlation matrix for 60s of MEG data, filtered in the gamma band using a Chebyshev band pass filter. The **top right** plot compares the first matrix with a similar matrix obtained by filtering for the delta band (key is the same as in Figure 17). Note that the correlation is clearly very strong implying that the correlations exist across a wide range of frequencies. The **second row** shows the comparison of the **top left** matrix with the Kuramoto correlation matrix at a range of $(k, \langle \tau \rangle)$ values, using only connected oscillator pairs. Note that as with the fMRI data and the connectivity matrix, at $(k=18, \langle \tau \rangle=11)$ the correlation is high and the error is low.

Although the parameter pair ($k=19, \langle \tau \rangle=11$) does appear to give a good match between the simulated correlations and the MEG data we need some way of estimating the confidence of this match. Here we employ two methods. Firstly we rerun the MEG projection 100 times adding uniformly distributed Cartesian jitter to the dipole locations. This is done for two magnitudes of jitter: the uniform distribution $[-5 \text{ mm}, +5 \text{ mm}]$ and the uniform distribution $[-10 \text{ mm}, +10 \text{ mm}]$. Secondly, we shuffle the rows/columns of the original MEG correlation matrix 10,000 times (equivalent to shuffling the points in space). Figure 22 shows the resulting three distributions of Pearson correlation values.

The conclusion to be drawn from Figure 22 is that the analysis is at most accurate to within 1 or 2cm. Also, we should note that most of the connected pairs in the network are physically close to each other, and thus are likely to have strong correlations according to the MEG projection, regardless of the underlying true activity.

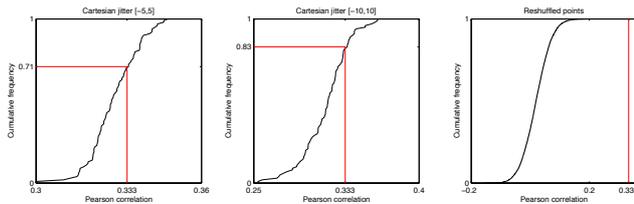


Figure 22: Cumulative distributions showing Pearson correlation values between disturbed MEG matrix and Kuramoto matrix for ($k=19, \langle \tau \rangle=11$). The undisturbed MEG matrix gives a Pearson correlation of 0.333, which is annotated here in red. See main text for full explanation.

5.3 Computation and the Kuramoto Model

We now move away from experimental results and examine the Kuramoto Model itself. Although we have done little so far to prove that the model is predictive of mesoscopic neural activity, it remains a reasonable claim.

One of the most complex Kuramoto models recently examined [31] used a network with eight ‘communities’ of 32 oscillators. The oscillators in each community were fully connected, with inter-community connections being both sparse and weak. The model used Equation 2 but with a fixed phase delay rather than the temporal offset, τ . The phase delay was the only parameter in the system; intra- and inter- community weights were fixed at two specific values.

The magnitude of the circular mean, $r_i(t)$, was calculated for each community and then fed into the calculation of three metrics: (a) the metastability, defined for each community as the variance of r over time, and for the whole system as the mean of this quantity over all communities; (b) the chimera-like quality of the system, defined for a given time as the variance of r over all communities, and for all time as the mean of this time-dependent quantity; (c) the coalition entropy, defined using Equation 5, where the value x indicates exactly which of the eight r values is greater than some fixed threshold (i.e. x takes 2^8 different values).

This system with its eight clusters is potentially similar to the mesoscopic brain (as we noted in Section 4.1). Thus, it is perhaps of interest that there is a critical range of the phase offset parameter which maximizes all three of the above metrics. This maximization implies that at the critical regime each community internally falls in and out of synchrony over time. Further, at any one time some communities are internally synchronized while others are not. And lastly, the selection of internally synchronized communities keeps changing. This highly dynamic behavior is considered to be a good model for a healthy brain [31].

A second recent Kuramoto implementation [11] was used to examine information flows during the onset of synchrony in a two-dimensional locally wired random network. Two analyses were performed on the time derivative of the oscillator phase: the active information storage was calculated, and the pairwise transfer entropy was calculated. Several conclusions were drawn from this analysis: (a) it appeared that information transfer only occurred in the first stage of the synchronization process (as measured by the magnitude of the circular mean); (b) active storage decreased and transfer entropy increased as coupling strength was increased; (c) active storage was negatively correlated with the degree of a node; (d) outgoing transfer entropy was positively correlated with the degree of a node.

In both of the above investigations the Kuramoto model is taken as being the single information carrying/computing mechanism in the system. However in Section 4.3 we noted that the synchronization of separate neural populations did not itself represent information, rather it (potentially) indicated the presence of a high-bandwidth connection between the regions. Given that individual

neurons can spike up to about 32 times a second, the bandwidth of a single axon is something like 5 bits per second if rate coding is used, perhaps more if temporal coding is used (c.f. analogue versus digital). Thus, even if the number of axons between two regions is small, and there is some degree of redundancy, the information transfer between the two regions may be relatively large. Further, note that the network used in this work includes only those inter-regional connections that are reasonably significant (due to the methodological limitations imposed in its construction).

This work is concerned with activity at the mesoscopic scale, so we must discard any notion of the true information flowing between regions and concentrate only on the gross properties of the regional activity. One question we can ask at this level is whether the system has multiple states and whether the evolution between states is periodic, totally chaotic, or follows a more complex pattern. This question relates to both the collation entropy used in [31] and the active storage used in [11]. Here however, rather than examine the temporal derivative of individual oscillators or the thresholded synchronization of small populations, we look at the phase difference between pairs of oscillators. The idea being that it is the phase difference between pairs of regions which is hypothesized to control the type of interaction in which they are engaged, i.e. excitatory, inhibitory, or null. Figure 23 shows an example of the type of data on which the analysis will be performed.

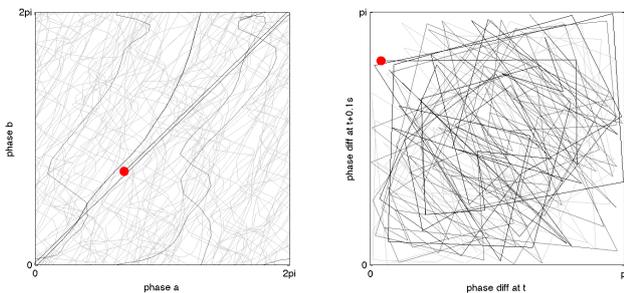


Figure 23: Both plots depict the same pair of oscillators from a Kuramoto simulation. The model described in Section 5.1 was used with $(k=18, \langle \tau \rangle=11)$, but run for only 23 seconds. The **left plot** shows the phase of one oscillator against the phase of the other, whereas the **right plot** shows how the phase difference changes every tenth of a second. In both plots the **red dot** shows where the two oscillators were initialized and the path starts off **black**, before turning **dark gray** and then **light gray**. Note that the left plot only shows the first 2.3 seconds of data, but the right plot shows the full 23 seconds.

All the relevant metrics introduced in Section 3.3 require that we know the probability of each particular sequence, $x_i, x_{i-1}, x_{i-2}, \dots$, where here x is going to represent the phase difference between a pair of oscillators, and we choose to downsample the data from 0.2ms to 0.1 seconds. A good estimate of the probability for a particular sequence can be obtained so long as the sequence occurs at least 10 times in the data [11]. However, given that the phase data is continuous we must specify what we mean by two sections of the data being the same.

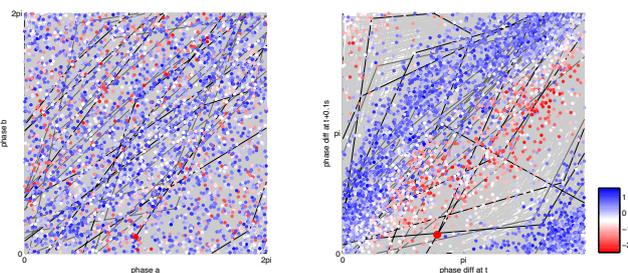


Figure 24: Both plots depict the same pair of oscillators from a Kuramoto simulation. The model described in Section 5.1 was used with $(k=18, \langle \tau \rangle=11)$, and run for 400 seconds. The **path** and **large red dot** have the same meaning as in Figure 23. The **colored dots** show the active information storage at each point.

Here we propose a kernel density algorithm that gives estimates for the probability of sequences of varying length, where the length is restricted by the requirement to have at least 10 occurrences of each sequence. For each point in time we calculate our metric with the maximum sized window, thus approximating the limit of the metric as the window size tends to infinity.

The algorithm was implemented in MATLAB de novo, but inspired by existing techniques [18, pp 257-261]. Here we give a brief explanation of how the algorithm works (see next page). Note that in the actual implementation used the outer most loop was vectorized.⁷

Figure 24 shows the active information storage for a single pair of oscillators in the model. Here we calculate active information storage for all 574 pairs of connected oscillators. Figure 25 shows that the variable-length window used does not introduce a bias.

⁷‘vectorisation’ is the MATLAB term for parallelisation with single instruction multiple data.

```

/** Algorithm for estimating  $\forall a \in T$ , the probability of the sequence  $\{x_{a-k_{backward}}, \dots, x_{a+k_{forward}}\}$  **/
h  $\leftarrow \frac{1}{3}\pi$  // resolution, i.e. range in which two values are equivalent
Nmin  $\leftarrow 10$  // minimum occurrences
T  $\leftarrow 4000$  // 10 samples for 400 seconds
for a = 1  $\rightarrow$  T do
  A  $\leftarrow \{t \in T : [x(t) - x(a)] \bmod 2\pi < h \text{ OR } [x(a) - x(t)] \bmod 2\pi < h\}$ 
  m  $\leftarrow T$ 
  k  $\leftarrow 0$ 
  while length(A) > Nmin do
    pa(k)  $\leftarrow \frac{1}{m}$  length(A)
    k  $\leftarrow k + 1$ 
    if k  $\leq k_{backward}$  then
      B  $\leftarrow A - k$ 
      B  $\leftarrow \{t \in B : [x(t) - x(a - k)] \bmod 2\pi < h \text{ OR } [x(t) - x(a - k)] \bmod 2\pi < h\}$ 
      A  $\leftarrow \{t \in A : t - k \in B\}$ 
      m  $\leftarrow m - 1$  // expanding window means one fewer samples to match against
    end if
    if k  $\leq k_{forward}$  then
      B  $\leftarrow A + k$ 
      B  $\leftarrow \{t \in B : [x(t) - x(a + k)] \bmod 2\pi < h \text{ OR } [x(t) - x(a + k)] \bmod 2\pi < h\}$ 
      A  $\leftarrow \{t \in A : t + k \in B\}$ 
      m  $\leftarrow m - 1$  // expanding window means one fewer samples to match against
    end if
  end while
end for

```

Low active information storage corresponds to oscillators deviating from an established pattern; high active information storage corresponds to a complex pattern being followed. Figure 26 shows the active storage values for all 574 pairs of oscillators. As with Figure 24, this data is somewhat difficult to interpret. Here we apply independent component analysis⁸ to examine whether there are groups of oscillator pairs that have a long term relationship in terms of the pattern of their phase difference changes. Figure 27 depicts the results of this analysis. It is still difficult to interpret whether there is anything to be learnt from the data, however here at least we see that there is some clustering in a manner similar to the connection matrix and the correlation matrix. The vertical and horizontal lines imply that some individual oscillators have a particularly strong influence on many other oscillators, however this may just be a result of significant heterogeneities in the connection matrix.

To test whether the values obtained are specific to the chosen parameter pair, the simulation was run a further two times at different points in parameter space. Figure 28 shows a simple comparison of these results, from which it appears that the ($k=18, \langle \tau \rangle=11$) simulation on average produces higher values of active information storage, potentially indicating a greater complexity of patterns.

6 Concluding Remarks

In this work we have explored the question of how the human brain operates at the mesoscopic scale. Section 4.3 was possibly the most important stage of this exploration as it was there that we discussed some of the theory of neural computation at the population level, before looking at how groups of populations might communicate across mesoscopic distances. Computational models, as introduced in Section 2, featured several times in this (and later) discussions as a tool to aid understanding of experimental results as well as a mechanism for introducing a hypothesis.

In Section 5 we attempted to show that the mesoscopic brain can be thought of as a network of weakly coupled oscillators. As with similar attempts in the past, this comparison proved difficult due to the inherent similarity between the experimentally obtained connectivity matrix and functional connectivity matrix. In addition we touched on - and attempted to put a bound on - the problems of spurious correlations arising in the source projections of MEG data. Finally, we examined the possibility of measuring patterns in the phase differences of pairs of the weakly coupled oscillators. This comparison proved difficult to interpret, but we may tentatively conclude that it captured the idea that at a critical point in parameter space the behavior of the model is more pattern-rich than elsewhere.

⁸Independent component analysis was applied using the freely available FastICA MATLAB toolbox.

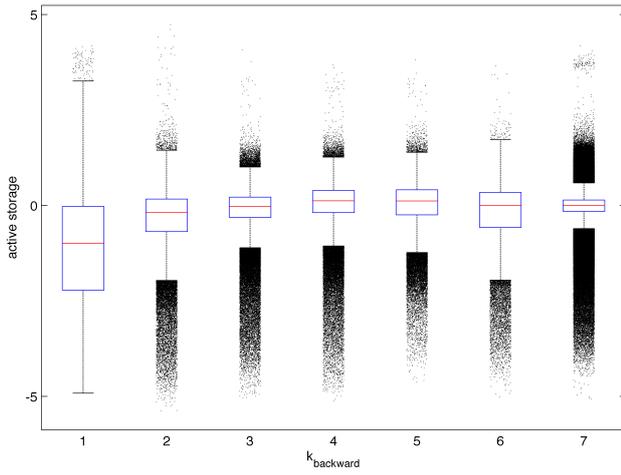


Figure 25: Box plot demonstrating the lack of obvious relationship between the window width and the active information storage. Note that 6% of points could not be assigned probabilities and are not shown. **Boxes** show inter quartile range, **whiskers** are a maximum of 1.5 times the interquartile range, and **individual points** are outliers.

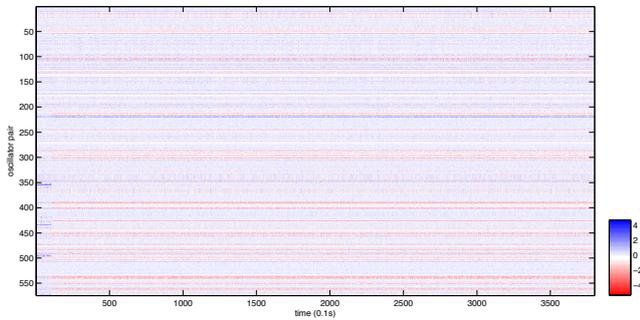


Figure 26: Plot showing the active information storage for all 574 pairs of connected oscillators.

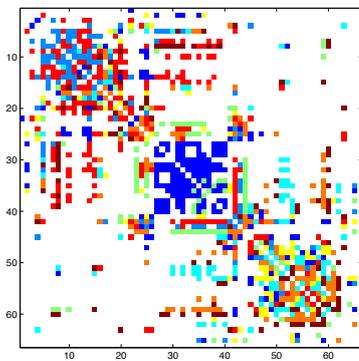


Figure 27: For each of the connected oscillator pairs this plot shows which of the (first eight) independent components it is most composed of.

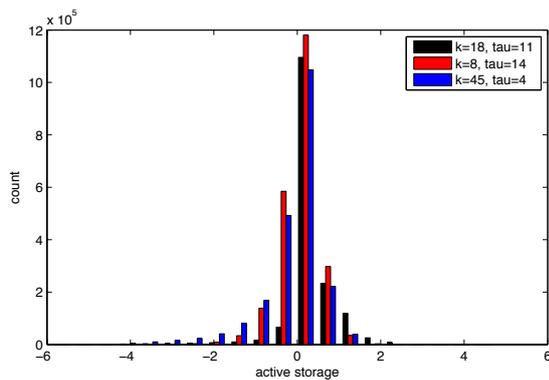


Figure 28: Distribution of the active information storage values calculated at each time point for the 574 oscillator pairs. **Legend** shows the three pairs of parameters used for the simulations. Note that approximately 6% of values in the first two simulations could not be assigned an active storage value, but the figure was 30% for the last simulation.

References

- [1] J.A. Acebrón, L.L. Bonilla, C.J.P. Vicente, F. Ritort, and R. Spigler. The kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of modern physics*, 77(1):137, 2005.
- [2] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters*, 59(4):381–384, 1987.
- [3] M. Bartos, I. Vida, and P. Jonas. Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews Neuroscience*, 8(1):45–56, 2007.
- [4] J.M. Beggs and N. Timme. Being critical of criticality in the brain. *Frontiers in Physiology*, 3, 2012.
- [5] M. Breakspear, S. Heitmann, and A. Daffertshofer. Generative models of cortical oscillations: neurobiological implications of the kuramoto model. *Frontiers in Human Neuroscience*, 4, 2010.
- [6] M.J. Brookes, M. Woolrich, H. Luckhoo, D. Price, J.R. Hale, M.C. Stephenson, G.R. Barnes, S.M. Smith, and P.G. Morris. Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *Proceedings of the National Academy of Sciences*, 108(40):16783–16788, 2011.
- [7] A. Buehlmann and G. Deco. Optimal information transfer in the cortex through synchronization. *PLoS computational biology*, 6(9):e1000934, 2010.
- [8] G. Buzsáki, C.A. Anastassiou, and C. Koch. The origin of extracellular fields and currentseeg, ecog, lfp and spikes. *Nature Reviews Neuroscience*, 13(6):407–420, 2012.
- [9] G. Buzsáki and A. Draguhn. Neuronal oscillations in cortical networks. *Science*, 304(5679):1926–1929, 2004.
- [10] J. Cabral, E. Hugues, O. Sporns, and G. Deco. Role of local network oscillations in resting-state functional connectivity. *NeuroImage*, 2011.
- [11] R.V. Ceguerra, J.T. Lizier, and A.Y. Zomaya. Information storage and transfer in the synchronization process in locally-connected networks. pages 54–61, 2011.
- [12] D.R. Chialvo. Emergent complex neural dynamics. *Nature Physics*, 6(10):744–750, 2010.
- [13] B.A. Cipra. An introduction to the ising model. *American Mathematical Monthly*, 94(10):937–959, 1987.
- [14] P. Fries. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in cognitive sciences*, 9(10):474–480, 2005.
- [15] P. Hagmann, M. Kurant, X. Gigandet, P. Thiran, V.J. Wedeen, R. Meuli, and J.P. Thiran. Mapping human whole-brain structural networks with diffusion mri. *PLoS One*, 2(7):e597, 2007.
- [16] CJ Honey, O. Sporns, L. Cammoun, X. Gigandet, J.P. Thiran, R. Meuli, and P. Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035, 2009.
- [17] M. Kaiser and C.C. Hilgetag. Nonoptimal component placement, but short processing paths, due to long-distance projections in neural systems. *PloS Computational Biology*, 2(7):e95, 2006.
- [18] H. Kantz, T. Schreiber, and R.S. Mackay. *Nonlinear time series analysis*, volume 2000. Cambridge university press Cambridge, 1997.
- [19] S. Katzner, I. Nauhaus, A. Benucci, V. Bonin, D.L. Ringach, and M. Carandini. Local origin of field potentials in visual cortex. *Neuron*, 61(1):35–41, 2009.
- [20] O. Kinouchi and M. Copelli. Optimal dynamical range of excitable networks at criticality. *Nature Physics*, 2(5):348–351, 2006.
- [21] C. Koch. *Biophysics of computation: information processing in single neurons*. Oxford University Press, USA, 1999.
- [22] R. Legenstein and W. Maass. Edge of chaos and prediction of computational performance for neural circuit models. *Neural Networks*, 20(3):323–334, 2007.
- [23] O. Levy, N.E. Ziv, and S. Marom. Enhancement of neural representation capacity by modular architecture in networks of cortical neurons. *European Journal of Neuroscience*, 2012.
- [24] J.T. Lizier, M. Prokopenko, and A.Y. Zomaya. Local measures of information storage in complex distributed computation. *Information Sciences*, 2012.
- [25] J.G. Milton. Neuronal avalanches, epileptic quakes and other transient forms of neurodynamics. *European Journal of Neuroscience*, 36(2):2156–2163, 2012.
- [26] E. Olbrich, P. Achermann, and T. Wennekers. The sleeping brain as a complex system. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1952):3697–3707, 2011.
- [27] R. Oostenveld, P. Fries, E. Maris, and J.M. Schoffelen. Fieldtrip: open source software for advanced analysis of meg, eeg, and invasive electrophysiological data. *Computational intelligence and neuroscience*, 2011:1, 2011.
- [28] N. Pržulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23(2):e177–e183, 2007.
- [29] E.T. Rolls, T.J. Webb, and G. Deco. Communication before coherence. *European Journal of Neuroscience*, 2012.
- [30] H. Sakaguchi and Y. Kuramoto. A soluble active rotator model showing phase transitions via mutual entrainment. *Prog. Theor. Phys*, 76(3):576–581, 1986.

- [31] M. Shanahan. Metastable chimera states in community-structured oscillator networks. *Arxiv preprint arXiv:0908.3881*, 2009.
- [32] W.L. Shew, H. Yang, T. Petermann, R. Roy, and D. Plenz. Neuronal avalanches imply maximum dynamic range in cortical networks at criticality. *The Journal of Neuroscience*, 29(49):15595–15600, 2009.
- [33] D. Sol. Revisiting the cognitive buffer hypothesis for the evolution of large brains. *Biology letters*, 5(1):130–133, 2009.
- [34] CJ Stam and ECW van Straaten. The organization of physiological brain networks. *Clinical Neurophysiology*, 2012.
- [35] S.H. Strogatz. *Sync: The emerging science of spontaneous order*. Hyperion, 2003.
- [36] E. Tulving. Organization of memory: Quo vadis. *The cognitive neurosciences*, pages 839–847, 1995.
- [37] A. M. Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, s2-42(1):230–265, 1937.
- [38] P.J. Uhlhaas, F. Roux, E. Rodriguez, A. Rotarska-Jagiela, and W. Singer. Neural synchrony and the development of cortical networks. *Trends in cognitive sciences*, 14(2):72–80, 2010.
- [39] B.D. Van Veen, W. Van Drongelen, M. Yuchtman, and A. Suzuki. Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *Biomedical Engineering, IEEE Transactions on*, 44(9):867–880, 1997.
- [40] J. Von Neumann, P.M. Churchland, and P.S. Churchland. *The computer and the brain*. Yale Univ Pr, 2000.
- [41] J. Vrba and SE Robinson. Linearly constrained minimum variance beamformers, synthetic aperture magnetometry, and music in meg applications. In *Signals, Systems and Computers, 2000. Conference Record of the Thirty-Fourth Asilomar Conference on*, volume 1, pages 313–317. IEEE, 2000.
- [42] RT Wicks, SC Chapman, and RO Dendy. Mutual information as a tool for identifying phase transitions in dynamical complex systems with limited data. *Physical Review E*, 75(5):051125, 2007.
- [43] S. Wolfram. Statistical mechanics of cellular automata. *Reviews of modern physics*, 55(3):601, 1983.
- [44] S. Wolfram and M. Gad-el Hak. A new kind of science. *Applied Mechanics Reviews*, 56:B18, 2003.
- [45] T. Womelsdorf, J.M. Schoffelen, R. Oostenveld, W. Singer, R. Desimone, A.K. Engel, and P. Fries. Modulation of neuronal interactions through neuronal synchronization. *Science*, 316(5831):1609–1612, 2007.

Acknowledgement. The author is grateful to Luc Berthouze, Simon Farmer, and Maria Botcharova for their direction technical assistant and support, and to Joana Cabral for providing her code.