# Chapter 4: The Turbulent Brain:

## Rhythms and Waves

**Professor Alain Goriely**

**18 March 2025**

The human brain is never at rest. Even in moments of stillness, in your thick skull beneath the skin and bones, a symphony of electrical rhythms pulses through its networks. These oscillations and waves of coordinated neuronal activity are the essential blocks upon which cognition, perception, and memory are built. From the slow undulations of deep sleep to the rapid bursts that mark sensory processing, brain rhythms dictate the way we think, learn, and experience the world. In this lecture, we will move beyond the brain geometry and its structural network that we have learned about in previous lectures to explore the brain's inner dynamics: how waves emerge, how they interact, and how they orchestrate the flow of information. The brain relies on oscillations to regulate every aspects of cognition such as attention, synchronising distant regions, and even shaping consciousness itself. What can these rhythms tell us about the nature and speed of thought? And what happens when they falter? By uncovering the principles of brain dynamics, we may come closer to answering one of neuroscience's most fundamental questions—from matter to mind, how does the brain, transform electric pulses into thought?

## 4.1 Electric man

The main discovery of brain waves is attributed to the German psychiatrist, Hans Berger (1873 –1941). The story goes that his interest with brain waves was sparked by a life-altering event, an incident that would shape his scientific pursuits for decades. As a young Prussian footman in 1893, Berger suffered a near-fatal accident when he was thrown from his horse and almost run down by a horse-drawn cart. To his astonishment, he discovered the same day that his sister, miles away, had sensed his distress and immediately sent him a telegram to check on him. Convinced that this was not a coincidence, Berger became obsessed with the idea that he had sent a pulse of of psychic energy to his sisters at the thought of dyin, long-distance *telepathy*. He later trained as a physician and spent the rest of his life seeking a scientific explanation, suspecting that electrical activity in the brain could generate magnetic disturbances capable of influencing others at a distance.

Berger's inquiries led him to a profound discovery. Through meticulous experimentation, he detected regular rhythmic oscillations in the brain's electrical activity ((Fig. 4.1). He observed that these waves, later named *alpha waves* (or *Berger waves*), were most prominent when subjects had their eyes closed but vanished upon visual engagement. This was the first glimpse

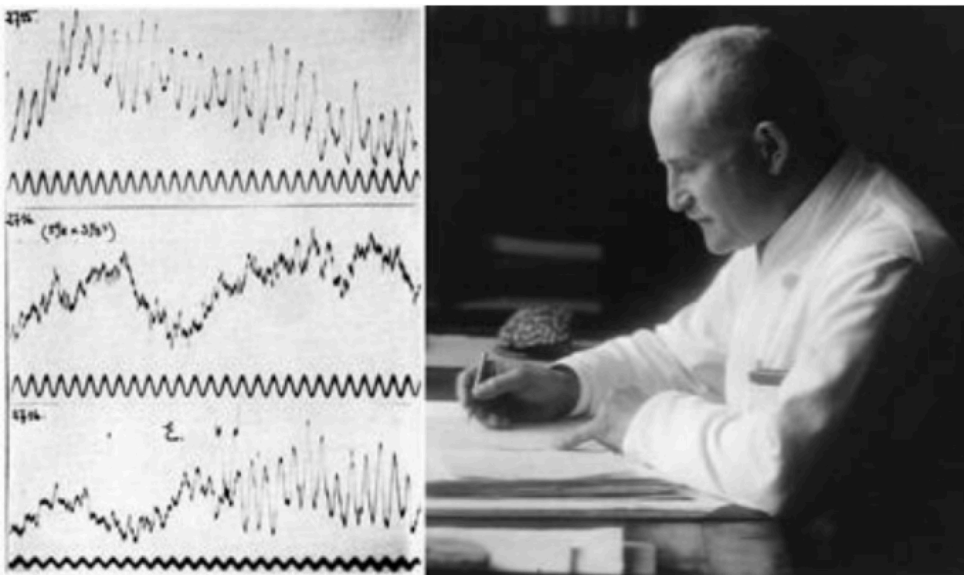

Figure 4.1: Hans Berger and the invention of EEG.

into brain rhythms, a discovery that would eventually revolutionise neuroscience. However, his discovery was met with great skepticism by the scientific community and mostly ignored by his colleagues in Germany, as many dismissed his findings as artefacts of flawed methodology or electrical interference. The dominant belief at the time was that the brain's activity was too chaotic to exhibit such organised oscillations. In particular, trying to disprove Berger's claims, the prominent physiologists Edgar Adrian attempted to replicate his experiments with more precise instrumentation. To their surprise, not only did they confirm Berger's findings, but Adrian went on to demonstrate that these rhythmic brain waves, particularly the alpha rhythm, were a fundamental and reproducible feature of brain activity. Adrian's endorsement in the 1930s ultimately validated Berger's discovery, and the method became known as *electroencephalography* (EEG), nowadays a basic tool in clinical neuroscience.

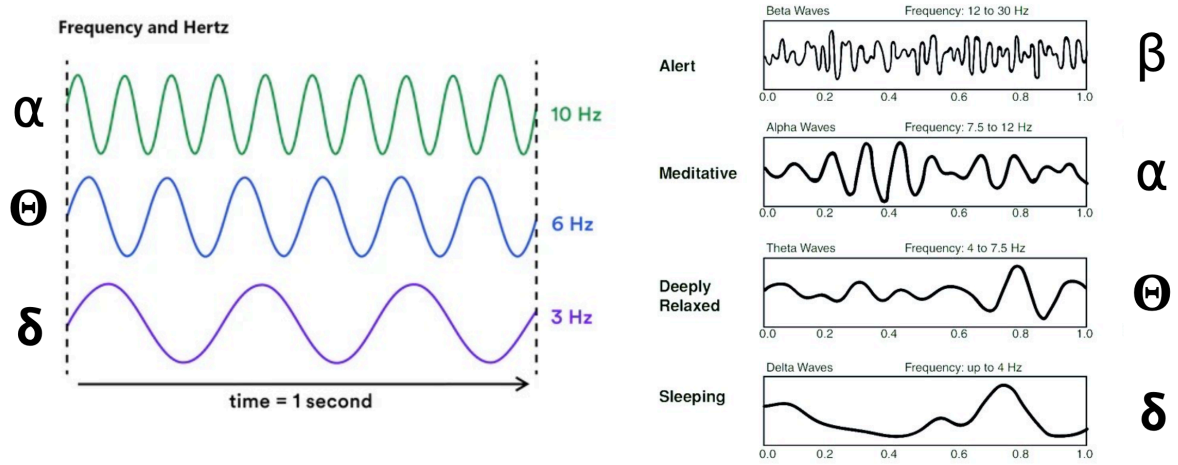


Figure 4.2: Depending on your mental state, brain waves have different frequencies from the slow delta waves in your sleep to beta waves associated with focus (or even higher, gamma waves).

We now categorise brain oscillations into five main frequency bands (Fig. 4.2), each associated

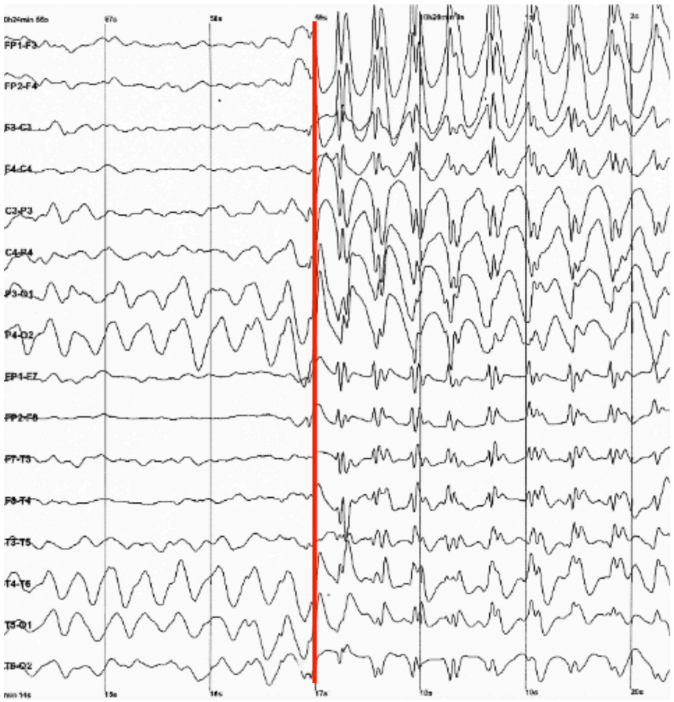

Figure 4.3: Generalized 3 Hz spike and wave discharges in a child with childhood absence epilepsy.

with different cognitive and physiological processes. At the slowest end of the spectrum are **delta waves** (0.5–4 Hz), the deep, powerful rhythms that dominate sleep and unconscious states. They facilitate rest, healing, and memory consolidation but can signal dysfunction when found in awake states. **Theta waves** (4–8 Hz), slightly faster, emerge during meditation, drowsiness, and moments of insight. They are thought to play a key role in creativity, memory retrieval, and emotional processing. Moving up the frequency ladder, **alpha waves** (8–12 Hz) are the brain's idle state, emerging when we are relaxed, yet alert. These oscillations are strongest when the eyes are closed and dissipate with sensory engagement, acting as a bridge between wakefulness and rest. Higher still are **beta waves** (12–30 Hz), which dominate active thinking, problem-solving, and motor control. These waves facilitate focused attention and cognitive efficiency but can also be linked to anxiety. At the top of the frequency spectrum, **gamma waves** (30–120 Hz) represent the fastest and most enigmatic oscillations and are associated with consciousness, deep learning, and the binding of sensory experiences into a unified perception.

Since Hans Berger's discovery of brain waves, EEG has been used as a non-invasive way to record the electrical activity of the brain, revealing how oscillations can become disrupted in pathological states. For instance, in epilepsy, seizures arise from abnormal, excessive synchronisation of neuronal firing, creating distinctive EEG signatures. These include spikes, sharp waves, and paroxysmal discharges, which indicate regions of the brain where hyperexcitability disrupts normal communication (Fig. 4.3). By capturing these abnormal rhythms, EEG allows clinicians to classify seizure types, locate epileptic foci, and guide treatment decisions.

The problem that we are faced with now is to explain how these waves can be generated by cellular activity and the interactions of many cells. We will start our exploration from micro to macro at the level of the cells.

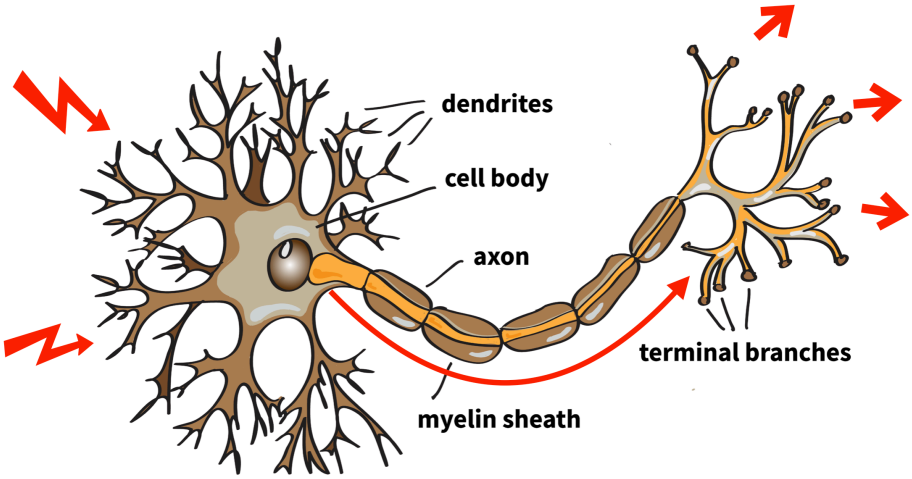


Figure 4.4: When a neuron receives a sufficiently strong signal from other neurons, an action potential is created that travels from the soma to the terminal branches of the axon.

## 4.2 Micro: Neurons and spikes

We learned in previous talks that the cells responsible for all cognitive functions are *neurons*, these are specialised cells that transmit electrical and chemical signals to process and relay information. More specifically, electrical signals is carried through the so-called *action potential*, an electrical signal that travels along a neuron's *axon* (the long filament protruding from the cell body). At rest the *membrane potential*, the difference between the exterior and interior of the neuron, is about -70mV (think of it as a very small battery). An action potential is then triggered when the cell's membrane potential reaches a critical threshold. This brief surge of voltage, is driven by the flow of sodium and potassium ions through voltage-gated channels in the axon membrane, allowing sodium ions to rush into the cell, causing depolarisation. If the membrane potential reaches the threshold (typically around -55mV), an *all-or-nothing* action potential is triggered. As the signal peaks, sodium channels close, and voltage-gated potassium channels open, allowing potassium ions to exit, leading to *repolarisation*. The neuron then becomes hyperpolarised and cannot be excited again, before returning to its resting state via the sodium-potassium pump, ready to fire again. This inability to fire and the all-or-nothing property enables neurons to transmit electrical impulses with great precision in a very noisy environment. In terms of information, one can think of a spike as one bit of information exchanged between different cells. Since the process involves the motion of ions in and out of the axon, it is actually very slow compared to an electric current in a wire, about a million times slower (electricity flows at about light speed whereas fast action potential are around 100 m/s, ). Yet, it is one of the fastest signalling processes in physiology.

### 4.2.1 The Hodgkin-Huxley model

The understanding of action potential is in great part due to the work of Alan Hodgkin and Andrew Huxley in the 1950s [4]. By studying the giant squid axon (*Loligo forbesi*), a uniquely large nerve fibre that allowed for precise experimental manipulations, they were able to record how voltage changes across the axon's membrane during an action potential. Doing so, they revealed the role of ion flow in neuronal signalling. Importantly for our story, their discovery was not just confined to experimental observation, but, in the great tradition of physiology, they identified the mechanisms responsible and demonstrated it by formulating a complete mathematical model describing how voltage-gated ion channels regulate neuronal excitability

[3]. This pioneering work is one of the greatest contributions to science in the 20th Century and earned them the 1963 Nobel Prize in Physiology or Medicine. It is the basis for all modern computational neuroscience, and inspired artificial intelligence.

More technically, the Hodgkin-Huxley model is a set of differential equations that describes the electrical properties of excitable cells. The model incorporates three major ionic species that contribute to the generation and propagation of action potentials and is formulated as an electric circuit.

The Hodgkin-Huxley model captures three distinct types of neuronal behaviour, which are also observed in real axons, depending on the strength of the external input (the current coming into the cell by interacting with other cells). These behaviours emerge from the way ion channels respond to voltage changes and determine whether an action potential is generated.

1. **For small inputs**, the neuron remains in its resting state and no spike occurs. This is because the external current is not strong enough to depolarise the membrane past the threshold potential, meaning that voltage-gated sodium channels do not open, and the neuron remains silent.

2. **For moderate inputs** but past a certain threshold, the neuron fires a single spike. If the input is strong enough to cross the threshold, sodium channels activate, leading to a rapid depolarisation and the generation of an action potential. However, after this single spike, potassium channels open to repolarise the neuron, and without further stimulation, the neuron returns to rest.

3. **For large inputs**, the neuron fires a train of spikes instead of just one. In this case, the external stimulus is strong enough that, after the first spike, the neuron does not return fully to rest before another action potential is triggered. This happens because sodium channels recover quickly, and the ongoing input keeps depolarising the neuron, leading to repetitive firing. The frequency of these spikes depends on the interaction between sodium influx and potassium efflux, which determines how fast the neuron resets between spikes.

These three behaviours—**no response, single spike, and repetitive firing**—are fundamental to how neurons process information and depend on an internal cell threshold that can be changed, making a cell more or less likely to react. In sensory neurons, weak stimuli may be ignored, moderate stimuli may trigger a brief response, and strong stimuli may lead to sustained signalling, allowing the brain to encode different intensities of input. The Hodgkin-Huxley equations not only predict these behaviours mathematically but also explain how neurons dynamically transition between them based on the properties of their ion channels.

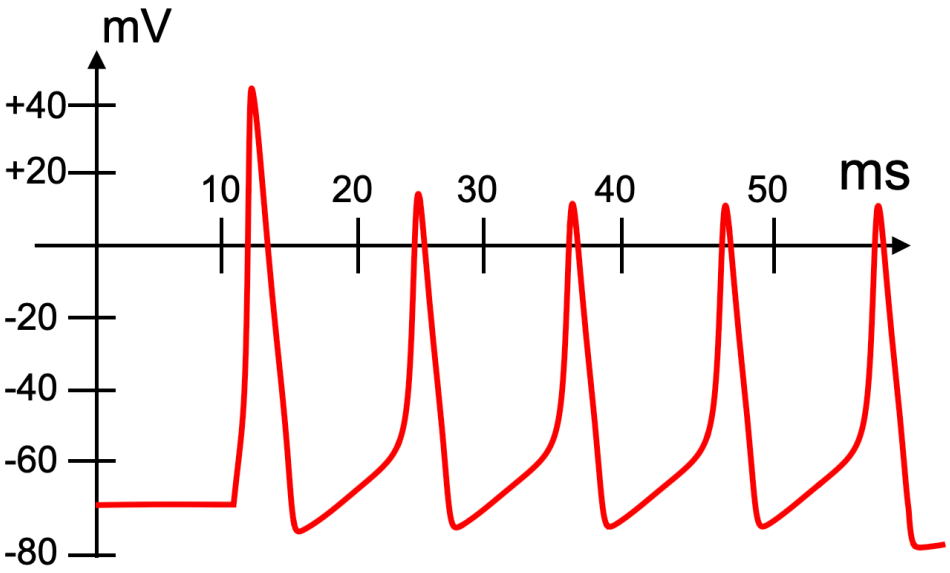


Figure 4.6: For larger values of the input, a train of spikes is generated.

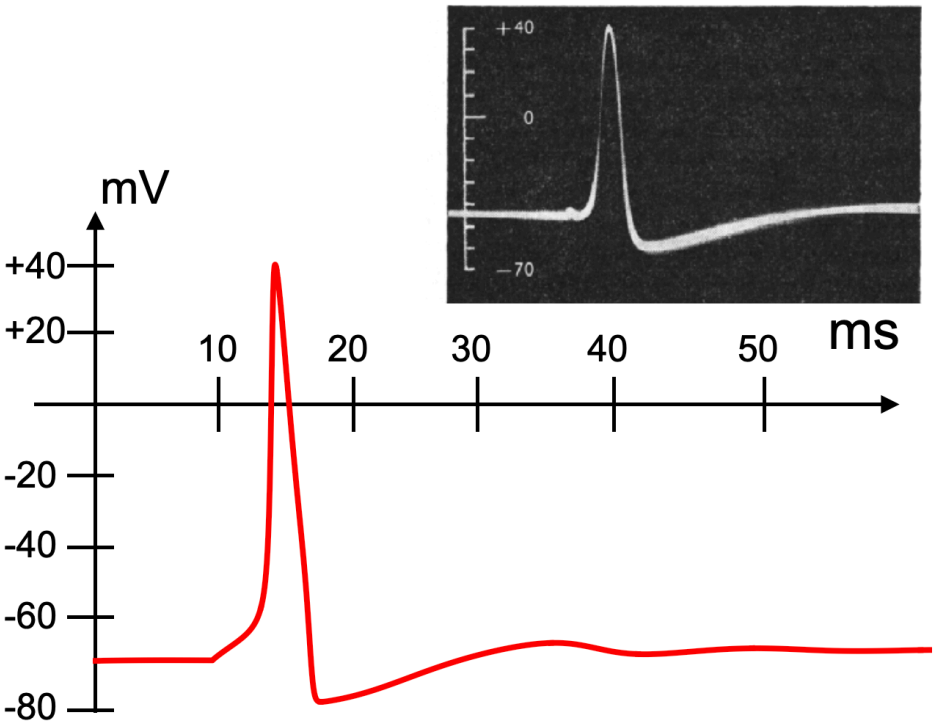


Figure 4.5: Past a threshold, the action potential travels along the axon. At a given point, the record shows a spike: a sudden increase of the action potential, followed by a return to the original value. Here, I show, the solution of the Hodgkin-Huxley model for the action potential as a function of time, compared to the original signal obtained by Hodgkin and Huxley.

The question is now how these isolated spikes can produce the brain oscillations that are observed. To do so, we now consider a small patch of tissue and expand our view to the meso-scale.

## 4.3 Meso: Waves

Neuronal activity emerges from the intricate interplay of *excitatory* and *inhibitory* neuron populations. These two types of neurons work in opposition, yet their balance is what enables the brain's dynamic and flexible computations. Excitatory neurons, the most common type of neurons, amplify signals and promote communication across circuits by releasing the neurotransmitter *glutamate*, which increases the likelihood that connected neurons will fire. In contrast, inhibitory neurons act as regulators, applying a crucial braking mechanism through the release of gamma-aminobutyric acid (GABA), which suppresses activity and prevents runaway excitation.

Brain waves result from an intricate feedback system between these two populations: when excitatory neurons drive bursts of firing, inhibitory neurons step in to restore balance, leading to the rhythmic patterns observed in EEG signals. In sensory processing, for example, excitatory neurons propagate signals from the environment, while inhibitory neurons refine and shape these responses, filtering out noise and ensuring only relevant information is passed along. In some circuits, such as those in the cortex and hippocampus, inhibitory neurons coordinate synchronised oscillations, crucial for cognitive functions like attention, learning, and memory.

Disruptions in this balance can have profound consequences. If inhibition is too weak, unchecked excitation can lead to pathological activity, as seen in *epileptic seizures*, where neurons fire excessively and synchronously. Conversely, excessive inhibition can dampen neural communication, contributing to disorders such as *schizophrenia* or *autism*, where information processing becomes inefficient. The careful coordination between these two neuronal populations is not just a matter of stability—it is the very foundation of all brain activity.

### 4.3.1 The Wilson-Cowan model

To understand how this feedback mechanism works at a mathematical level, we can model the interaction of these two populations of neurons by the Wilson-Cowan model first proposed in 1972 [8, 2]. In this model, instead of focusing on single neurons, we look at the average activity of the two populations of neurons: Excitatory neurons increase activity by stimulating themselves and their neighbours, and inhibitory neurons that act as a brake, suppressing activity to prevent uncontrolled firing. When excitatory neurons receive enough input, they increase their firing rate, which can then excite more neurons. However, as activity rises, inhibitory neurons also become more active, eventually suppressing excitatory activity. This interplay can lead to different patterns of activity, including steady states, rhythmic oscillations, or bursts of activity, similar to the patterns seen in EEG brain waves. Explicitly, the model, sketched in Fig. 4.7 consists in two equations for the *average activity levels* $E$ and $I$ of the *excitatory* and *inhibitory* neuron populations as a function of time

$$\frac{\mathrm{d}E}{\mathrm{d}t} = -E + S_E\left(w_{EE}E - w_{EI}I + P_E\right) \tag{4.1}$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = -I + S_I\left(w_{IE}E - w_{II}I + P_I\right), \tag{4.2}$$

where

- $\frac{\mathrm{d}E}{\mathrm{d}t}$ and $\frac{\mathrm{d}I}{\mathrm{d}t}$ represent the excitatory and inhibitory activity change over time.
- The first terms $-E$ and $-I$account for *natural decay* of neural activity over time. If no external input or interaction occurs, activity will slowly decrease.
- $S_E(\cdot)$ and $S_I(\cdot)$: These are *activation functions* (also called "sigmoid functions") that determine how the total input to each population translates into neural activity. They intro-

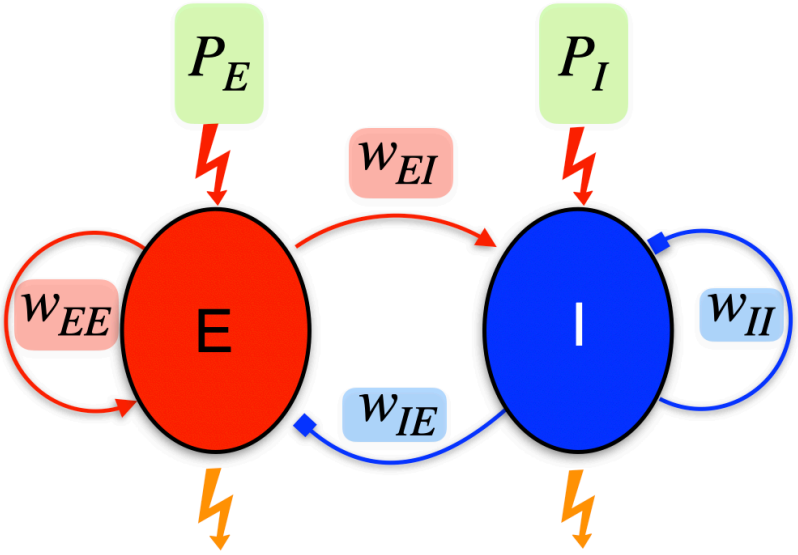


Figure 4.7: In the Wilson-Cowan model, we consider two populations of neurons: the excitatory and inhibitory populations.

duce *non-linearity*, meaning that neurons do not simply add up their inputs but instead respond in a way that saturates at high input levels.

- The term $w_{EE}E$ represents *excitatory self-feedback*, meaning that excitatory neurons stimulate each other, increasing $E$ when $E$ is high.

- The term $-w_{EI}I$ is the *inhibitory effect on excitatory neurons*. A strong inhibitory population suppresses the excitatory activity.

- The term $w_{IE}E$represents the *excitatory drive on inhibitory neurons*. When excitatory neurons are highly active, they increase the activity of inhibitory neurons.

- The term $-w_{II}I$: This term describes *inhibitory self-regulation*, meaning that inhibitory neurons also suppress themselves to prevent excessive inhibition.

- $P_E$ and $P_I$: These are *external inputs* to the excitatory and inhibitory populations. They can represent sensory stimuli, inputs from other brain areas, or artificial stimulation in an experiment.

Intuitively, If $P_E$ (excitatory input) is small, $E$ may not reach the threshold to fire, leading to *no activity*. If $P_E$ is *moderate*, just like in the Hodgkin-Huxley model, $E$ crosses the threshold and produces a *single spike* before inhibition suppresses it. If $P_E$ is *large or sustained*, the excitation keeps activating inhibition in a cycle, leading to *oscillations* or a *train of spikes*, as seen in EEG rhythms.

Mathematically, it is very easy to simulate numerically these equations and plot the solutions. We find that for suitable values of the parameters, we recover the oscillatory behaviour observed in EEG as shown in Fig. 4.7

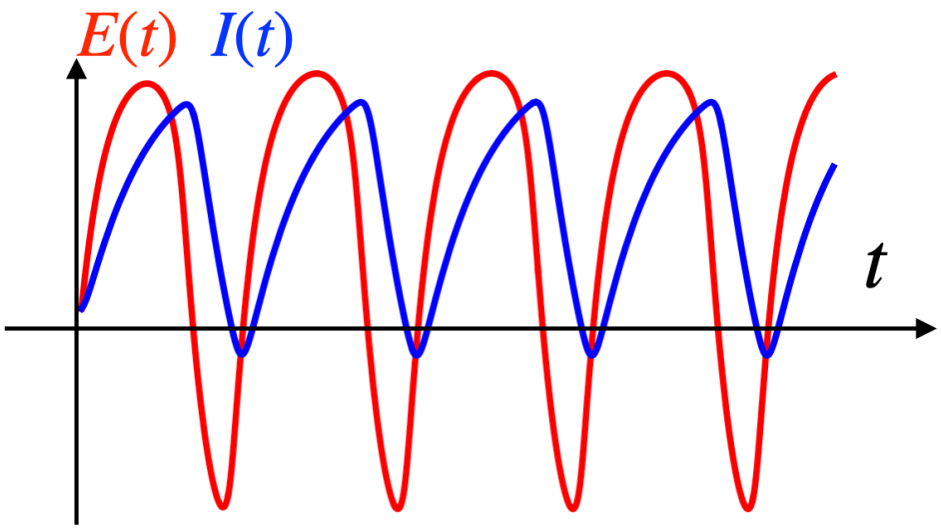


Figure 4.8: For certain values of the parameters in the model, oscillations are generated, as observed in EEG.

We note that the sigmoid function that is found in this model is also fundamental in all brain models because it introduces non-linearity into neural dynamics. In our model, the sigmoid function determines how the total input to a neural population translates into activity, mimicking the way real neurons respond to stimuli. Instead of reacting linearly to input, neurons exhibit threshold behaviour: when the input is too low, they remain inactive; when it is moderate, they respond proportionally; and when it is too high, they saturate, preventing runaway activity. This smooth transition from inaction to full activation is essential for maintaining stable and controlled neural responses, both in biological brains and in mathematical models of neural circuits.

Inspired by neuroscience, the sigmoid function played a crucial role in the development of artificial neural networks. In the 1980s, John Hopfield used it in his model of associative memory, where networks of neurons interact through weighted connections to store and retrieve patterns. Around the same time, Geoffrey Hinton and his collaborators incorporated the sigmoid function into *back-propagation*, the key algorithm for training artificial neural networks. For their contributions, they received the 2024 Nobel Prize in physics. Although modern deep learning has largely moved toward other activation functions, the sigmoid function remains historically significant. It was one of the first biologically inspired activation functions to bridge the gap between neuroscience and artificial intelligence, shaping the foundations of modern neural networks.

Now that we understand how a small patch of tissue behave, we can link all these regions in a brain at the macro-scale.

## 4.4 Macro: Networks

To build a model of the brain, we can extend the Wilson-Cowan model from a single population to a network of interacting populations. Instead of modelling just one group of excitatory and inhibitory neurons, we assign a Wilson-Cowan system to each node in a network, where each node represents a distinct brain region or cortical column as shown in Fig. 4.9. These nodes are then connected to each other, allowing activity to propagate between them, mimicking real neural communication.

Mathematically, each node in the network follows its own set of Wilson-Cowan equations, describing the local interactions between excitatory and inhibitory neurons. However, these equations are now coupled to those of other nodes through connection weights, which determine how strongly activity at one node influences another. If a node represents a brain region, these weights can be derived from connectome data, which maps the physical connectivity between different areas of the brain.

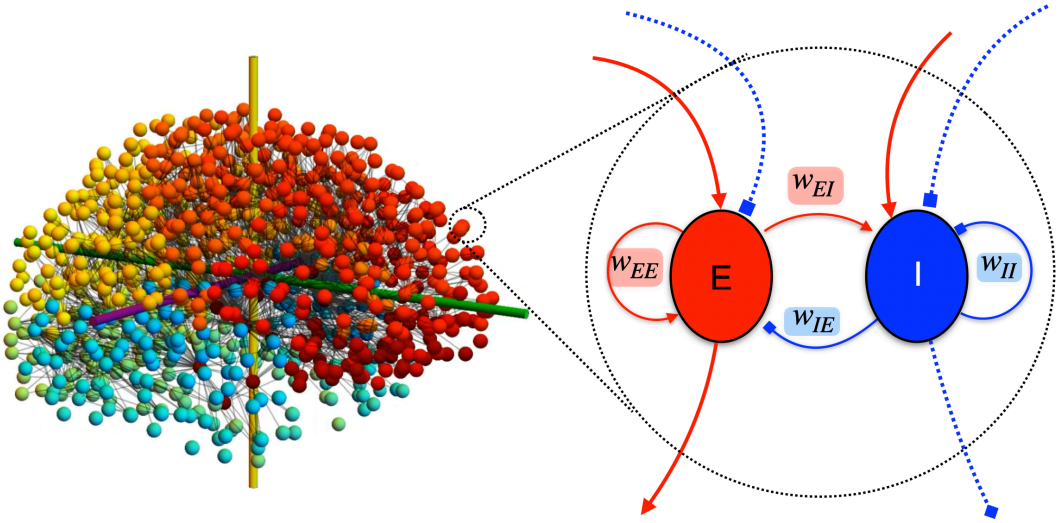


Figure 4.9: To understand brain dynamics, the activity at each node in a network is represented by two local populations of the Wilson-Cowan model with inputs and outputs from other nodes.

The resulting system is a network of coupled neural populations, where activity can spread dynamically across nodes, much like how signals travel through the brain. This allows us to study large-scale brain dynamics, including oscillations, wave propagation, and synchronization patterns observed in EEG and fMRI data. By adjusting the strength of connections, the external inputs, or the properties of individual nodes, we can simulate various brain states, from resting activity to pathological conditions like epilepsy. This approach provides a powerful framework for understanding how local neural circuits interact to create global brain function.

## 4.5 Be BOLD

While EEG has very good temporal resolution, it has poor spatial resolution and can only record cortical activity. An alternative is to use the **BOLD (Blood Oxygen Level Dependent)** signal, the primary measure used in **functional magnetic resonance imaging (fMRI)** to study brain activity. It does not directly record neural activity but instead captures changes in blood flow and oxygenation, which are closely linked to neuronal function. When a brain region becomes active, its neurons need energy and therefore consume more oxygen, triggering a hemodynamic response where nearby blood vessels dilate to supply fresh, oxygen-rich blood, as shown in Fig. 4.10. This process alters the ratio of oxygenated to deoxygenated hemoglobin, which affects how magnetic fields interact with tissue and can be detected by an MRI scanner.

Although BOLD is an indirect measure of brain function, it has proven incredibly powerful for mapping brain connectivity. To build a functional connectome, we look at correlated BOLD activity across different brain regions. If two regions show synchronised fluctuations in their BOLD signals over time, they are considered functionally connected, meaning they likely work together in cognitive processes. This functional connectivity is mapped across the entire brain,

creating a network model where nodes represent brain regions and edges represent their statistical correlation.

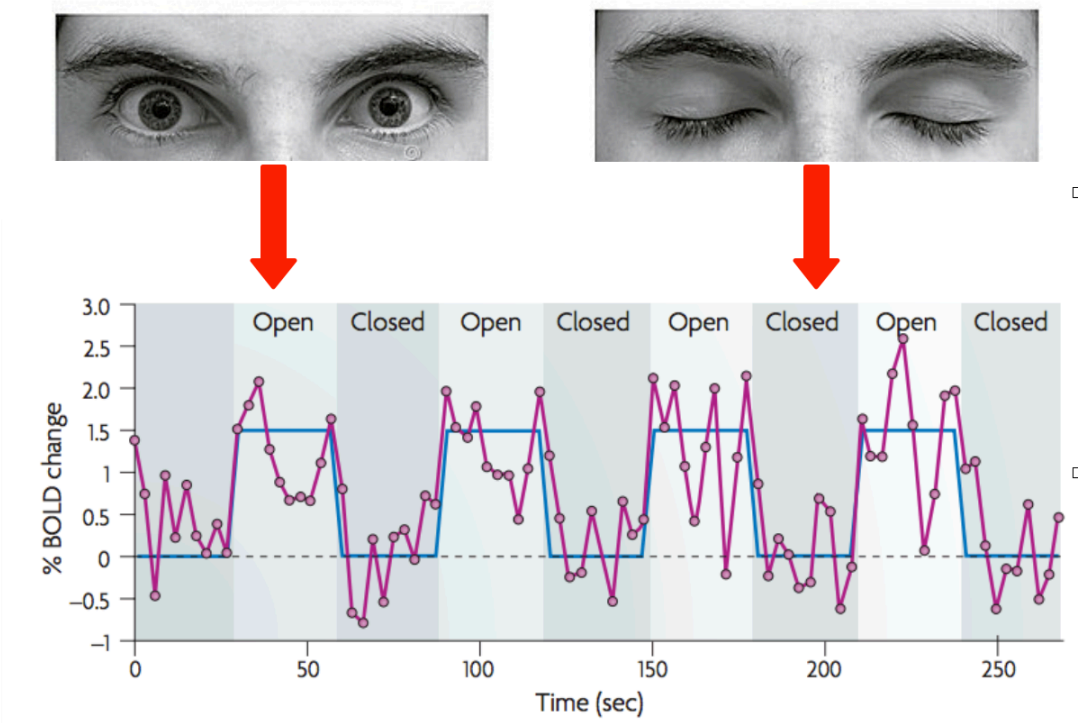


Figure 4.10: Measuring the BOLD signal at different places of the brain (here in the visual cortex) gives information about particular brain activity and their effects on different regions.

The correlations used in BOLD-based functional connectivity is a statistical measure that quantifies how similar the activity patterns of two brain regions are over time. Since the BOLD signal fluctuates in response to neuronal activity, we can analyse these fluctuations across different brain regions to determine how strongly they are functionally connected.

To compute **functional connectivity**, the most common approach is Pearson correlation (see first lecture), which measures the linear relationship between the BOLD signals of two regions. If the BOLD activity in two regions rises and falls together over time, they have a high positive correlation (close to +1), as in Fig. 4.11. If one region's activity increases while the other decreases, they have a negative correlation (close to -1). If there is no consistent relationship between their fluctuations, the correlation is near zero, suggesting little or no functional interaction.

Mathematically, the correlation coefficient $r$ between two time series $X(t)$ and $Y(t)$ is calculated as:

$$r = \frac{\sum(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum(X_i - \bar{X})^2}\sqrt{\sum(Y_i - \bar{Y})^2}} \tag{4.3}$$

where $X_i$ and $Y_i$ represent the BOLD signal values at a given time point, and $\bar{X}$ and $\bar{Y}$ are their mean values.

Once correlations are computed for all pairs of brain regions, we can construct a functional connectivity matrix, where each entry represents the correlation between two regions. This matrix is often visualised as a functional connectome, a network where brain regions (nodes) are connected by edges weighted by their correlation strength.

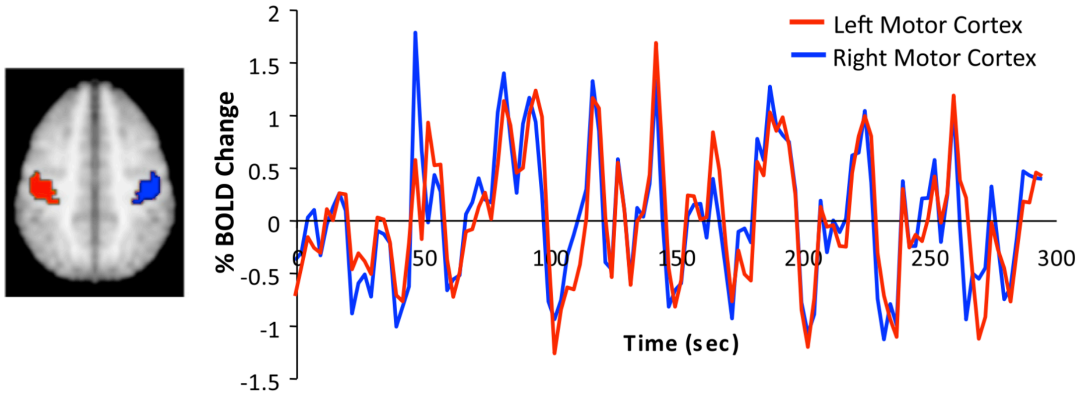


Figure 4.11: Comparing BOLD signal at different places in the brain tells us how different regions are linked functionally.

### 4.5.1 The functional connectome

Whereas structural connectomes are hard-wired networks in the brain, there is another type of networks that encode how different parts of the brain interact with each other. A **functional connectome** based on how different parts of the brain are activated. If they are synchronised and activated in the same way, then they are connected. Hence, functional networks capture the dynamic interactions between brain regions based on their activity. For instance, when you are solving a math problem or daydreaming, different sets of brain areas light up and synchronise, forming temporary functional networks tailored to the task at hand. These networks have revolutionised our understanding of the brain and are now routinely used in neuroscience.

One of the most exciting discoveries about functional networks is the brain's **default mode network**. It is a large-scale brain network that becomes most active when the mind is at rest, engaged in self-referential thinking, daydreaming, autobiographical memory, and introspection. The default mode network deactivates during externally focused tasks, shifting control to attention and executive networks. The default mode network is not an isolated system but dynamically interacts with other brain networks. When attention shifts away from introspection or self-referential thinking, the default mode network is deactivated: it shows anti-correlation with task-positive networks such as the **dorsal attention network** , which governs goal-directed behaviour and sensory processing. This interplay allows the brain to efficiently allocate resources, suppressing internal reflection when external engagement is required. Similarly, **the salience network**, which detects relevant stimuli and modulates attention, acts as a switch between the default mode network and externally focused networks, determining whether the brain remains in a state of rest or shifts toward active cognitive engagement. Understanding these cross-network interactions offers deeper insight into how the brain balances spontaneous thought, attention, and executive control to support complex cognition.

Functional networks reveal the brain's flexibility and adaptability, helping us understand how it reconfigures itself for everything from deep concentration to creative thinking. They are the key to unlocking how the brain's macro network team up together to create the mind.

## Epilogue

While we have explored the intricate dynamics of the fast oscillations that govern thought, perception, and action, there is a far slower and more insidious form of brain dynamics that we have yet to consider. Beyond the fleeting pulses of neural activity, another process unfolds

over months, years, and decades—a gradual, relentless alteration of the brain's structure and function. These are the changes brought about by neurodegenerative diseases, the cruelest of all brain disorders, in which the architecture of the mind slowly unravels. Unlike the rapid spikes and waves of healthy cognition, these pathological processes progress imperceptibly at first, dismantling neural circuits in ways that are often irreversible. Understanding these slow degradations is one of the greatest challenges of neuroscience, and in our next lecture, we will see how mathematics can help understand the dynamics of neurodegeneration.

## Further reading

- *Rhythms of the Brain* by György Buzsáki (2006) is an exploration of brain oscillations and their role in cognition, perception, and behaviour [1].
- *Networks of the Brain* (2016) by Olaf Sporns, one of the leaders of the field is a technical but very readable book [6].
- *Sync: How Order Emerges from Chaos in the Universe, Nature, and Daily Life* by Steven Strogatz (2003). Although not strictly about the brain, this book explains the mathematics behind synchronization, a key feature of brain waves and oscillatory neural activity [7]..
- *The entangled brain: How perception, cognition, and emotion are woven together* (2022) by Luiz Pessoa is well written and informative presenting a more in-depth analysis of brain networks than most books [5].

## References and Further Reading


[1] György Buzsáki. *Rhythms of the Brain*. Oxford university press, 2006.

[2] Alain Destexhe and Terrence J Sejnowski. The wilson–cowan model, 36 years later. *Biological cybernetics*, 101(1):1–2, 2009.

[3] Michael Häusser. The hodgkin-huxley theory of the action potential. *Nature neuroscience*, 3(11):1165–1165, 2000.

[4] Alan L. Hodgkin and Andrew F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, 1952.

[5] Luiz Pessoa. *The entangled brain: How perception, cognition, and emotion are woven together*. MIT Press, 2022.

[6] Olaf Sporns. *Networks of the Brain*. MIT press, 2016.

[7] Steven Strogatz. Sync: The emerging science of spontaneous order. 2004.

[8] Hugh R Wilson and Jack D Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12(1):1–24, 1972.