

Review

The Thermodynamics of Mind

Morten L. Kringelbach ^{1,2,3,4,*}, Yonatan Sanz Perl^{4,5,6}, and Gustavo Deco^{4,5,7,*}

To not only survive, but also thrive, the brain must efficiently orchestrate distributed computation across space and time. This requires hierarchical organisation facilitating fast information transfer and processing at the lowest possible metabolic cost. Quantifying brain hierarchy is difficult but can be estimated from the asymmetry of information flow. Thermodynamics has successfully characterised hierarchy in many other complex systems. Here, we propose the ‘Thermodynamics of Mind’ framework as a natural way to quantify hierarchical brain orchestration and its underlying mechanisms. This has already provided novel insights into the orchestration of hierarchy in brain states including movie watching, where the hierarchy of the brain is flatter than during rest. Overall, this framework holds great promise for revealing the orchestration of cognition.

Understanding the orchestration of brain dynamics

[Thermodynamics] is the only physical theory of universal content, which I am convinced will never be overthrown –

[Albert Einstein]

There is a conundrum at the heart of neuroscience, namely how billions of relatively slow neurons can carry out the computations needed for the flexible, time-critical behaviour needed for survival. This slowness arises from the electrical signals of a neuron being converted to a chemical signal at the synaptic junction before being converted back to an electrical signal [1]. The speed of information transfer between neurons is typically on the order of ~10–20 ms [2,3], which is several orders of magnitude slower than that found in computers. Yet, the brain is often better at solving difficult problems compared with a computer. The answer to this conundrum lies in the hierarchical architecture of the mammalian brain, which allows for the computation of sensory input followed by higher-level computation in nested recursive circuits at various spatio-temporal levels [4–6]. Although information flow is primarily shaped by brain anatomy, not unlike dynamic traffic flow on an existing road network, the sculpting effects of neurotransmission [7–9] provide additional flexibility. Ultimately, hierarchical brain processing allows information to be segregated and integrated as needed and facilitates the execution of the time-critical computations needed for survival [10–12].

Nevertheless, what is needed is a much deeper understanding of how this hierarchical processing is driven, or orchestrated, in different brain states, such as wakefulness, sleep, and anaesthesia [11,13,14]. However, even defining brain states is difficult and not commonly agreed upon [15]. Brain states clearly differ from each other in terms of their continuously evolving dynamics of whole-brain networks characterised by condition-dependent self-organisation in stable, semistable, and transient arrangements. There are influential theories trying to explain brain function, such as hierarchical models of predictive coding [16], hierarchical core-periphery principles [17–20], and the free-energy principles of the Bayesian brain [21,22]. However, these theories have not agreed on a common definition of a brain state and, while they

Highlights

To survive, the brain must perform fast, efficient distributed computation.

The challenge is to discover the orchestration of the brain hierarchy over space and time.

Thermodynamics is a natural framework to quantify hierarchy in any system.

The key concepts of irreversibility and the ‘arrow of time’ can reveal the asymmetry of brain information flow.

The ‘Thermodynamics of Mind’ framework has provided novel insights into the changing hierarchy of brain states.

Combined with whole-brain modelling, thermodynamics shows great promise for revealing causal insights into the orchestration of cognition.

¹Centre for Eudaimonia and Human Flourishing, Linacre College, University of Oxford, Oxford, UK

²Department of Psychiatry, University of Oxford, Oxford, UK

³Center for Music in the Brain, Department of Clinical Medicine, Aarhus University, Aarhus, Denmark

⁴International Centre for Flourishing, Universities of Oxford, Aarhus, and Pompeu Fabra, Oxford, UK

⁵Center for Brain and Cognition, Computational Neuroscience Group, Department of Information and Communication Technologies, Universitat Pompeu Fabra, Roc Boronat 138, Barcelona, Spain

⁶Department of Physics, University of Buenos Aires, Buenos Aires, Argentina

⁷Institució Catalana de la Recerca i Estudis Avançats (ICREA), Passeig Lluís Companys 23, Barcelona, Spain

*Correspondence:
morten.kringelbach@psych.ox.ac.uk
(M.L. Kringelbach) and
gustavo.deco@upf.edu (G. Deco).



acknowledge the hierarchical nature of brain processing, these theories have not been successful in characterising the causal mechanisms of whole-brain dynamics, in large part due to the complexity of the brain.

Moving beyond existing theories, a general theory of the brain needs to directly quantify brain states in terms of hierarchy at the level of whole-brain dynamics and to offer novel insights into the underlying mechanisms of brain states. The ‘Thermodynamics of Mind’ framework is a step toward a general theory quantifying hierarchy in brain states by using thermodynamics principles to quantify the underlying asymmetry of hierarchy. As an example, think of how egalitarian organisations have a flat, symmetrical structure, while many enterprises have a pyramidal structure with one or more leaders at the top, who delegate to many lower levels of management. Establishing the hierarchy in these organisations relies on establishing the asymmetry of information flow in the organisation.

These hierarchical principles also hold for the brain, where the asymmetry of the information flow determines the functional hierarchy. Thermodynamics provides a general framework for studying hierarchy in physical systems, including systems biology [23]. Specifically, this is provided by the specific branch of nonequilibrium thermodynamics that quantifies the asymmetry of information flow by estimating the reversibility and non-reversibility (irreversibility) over time of the underlying processes. To understand this key concept of irreversibility (‘arrow of time’; Box 1), think of watching a film of a glass being shattered by a bullet. This is a clear example of an irreversible process, where the glass goes from order to disorder. By contrast, when watching the same film sequence in reverse, we immediately recognise that, in the real physical world, it is not possible for the glass to come back together since this violates the second law of thermodynamics by going from disorder to order. The impossibility and, therefore, irreversibility of the events are abundantly clear. Applying the thermodynamical principle of irreversibility to the complexity of the brain allows for robust estimation of functional brain hierarchy. This is achieved by using irreversibility to quantify the asymmetry of information flow between all brain regions. More fundamentally, this can also provide insight into the mechanisms generating this asymmetry.

Here, we describe recent progress in using thermodynamics to describe the fundamental mechanisms underlying the orchestration of hierarchical brain dynamics. First, we briefly describe how scientists have tried to establish the functional hierarchy in the brain. We demonstrate how thermodynamics can directly provide both quantification and underlying mechanisms of hierarchy through asymmetry. The Thermodynamics of Mind framework can explain how brain orchestration of hierarchy ensures survival. We provide three examples of how this causal framework has led to promising new insights in neuroscience and psychology, in terms of (i) the orchestration of information flow in brain states and cognition; (ii) what happens during movie watching; and (iii) how, despite its relative slowness, the brain can process information fast enough to survive. Finally, we explore some of the many fertile avenues of research arising from the framework.

Understanding brain hierarchy

Hierarchy can be precisely described mathematically using order theory and, in particular, **partially ordered set (poset)**; see Glossary), which formalises the ordering, sequencing, or arrangement of the elements of a set [24]. Hierarchy is an organising principle in all living systems [25]. This can be appreciated by taking a thermodynamic approach to modelling biological systems as physical systems, where, in the most general abstraction, they are thermodynamic open systems showing self-organised behaviour. The set-subset relations between dissipative structures can be characterised by a hierarchy across spatiotemporal scales.

Glossary

Fluctuation–dissipation theorem

(FDT): a central theorem in statistical thermodynamics for predicting the nonequilibrium fluctuations of a system, such as the irreversible dissipation of energy into heat from its reversible fluctuations in thermal equilibrium.

Generative mechanisms: generative mechanisms underlying the temporal evolution of a system can be determined by building a model of the system and investigating the causal influence of manipulating the model elements (also see ‘whole-brain model’).

Granger causality: method developed by econometrician Clive Granger for testing how useful one time series is for forecasting another time series. Importantly, despite its name, this method does not directly measure causality but merely the degree of temporal correlation.

Partially ordered set (poset): the mathematical definition of hierarchy relies on a partially ordered set (poset) coming from the order theory branch of mathematics. This formalises the ordering, sequencing, or arrangement of the elements of the entire poset.

Power law: special mathematical relationship between two quantities in which one quantity varies as a power of the other. The power law distribution arises when extreme events occur with low probability, such as how most people in a social network only have a couple of hundred contacts, while some influencers may have millions. The power law often indicates that a system such as the brain could be scale free and operate in a critical state of self-organised criticality, which makes the system highly robust to random failures, but vulnerable to attacks.

Whole-brain model: powerful tool for modelling brain dynamics from whole-brain neuroimaging techniques, such as fMRI or magnetoencephalography. In its simplest form, the whole-brain model is constructed using the anatomical connectivity of a reduced set of typically hundreds of anatomically defined brain regions. Each anatomically linked region contains a model of the local dynamics, and the model is fitted to the neuroimaging time series by simply scaling the global connectivity. The elements of such an *in silico* model of brain dynamics can then be exhaustively probed, and the underlying causal mechanisms revealed.

Box 1. The arrow of time in thermodynamics

The nonequilibrium thermodynamic principle of the arrow of time can be illustrated with sequences from two films. The first sequence of images is a classic example of an equilibrium system and is taken from a film of colliding billiard balls (Figure 1A). When watching this film both forward and backward, the sequences are almost identical, and it is almost impossible to distinguish the direction of the arrow of time for each of the films. In thermodynamical terms, this is because the process does not produce entropy and creates a reversible process. By contrast, the sequence of images of a movie of glass being shattered (Figure 1B) is a strong example of a nonequilibrium system and of how irreversible changes lead to an increase in production entropy. Equally, when watching the reversed film sequence of the glass being shattered, it is intuitively clear that the second law of thermodynamics means that a glass cannot spontaneously come back together, that is, the transition from disorder to order is impossible. This establishes a clear arrow of time, where the forward and backward unfolding of events are distinguishable.

More generally, irreversibility is closely linked with production entropy [71] as illustrated in Figure 1C, which shows a nonequilibrium system with two states A and B, and the associated trajectories evolving during forward ($A \rightarrow B$, black arrow) and backward ($B \rightarrow A$, red arrow) processes [72]. Both the forward and backward trajectories can be thought of as corresponding to the movies shown in the previous panels, but each with a different arrow of time. By contrast, the time reversal of the backward trajectory (red stippled arrow) can be imagined as the movie of the backward trajectory played forward in time. If the forward and time reversals of the backward trajectories are different, this corresponds to non-reversibility of the process.

Finally, Figure 1D shows the second law of thermodynamics stated in terms of the concept of production entropy [23,36]. If the production entropy, H_p , is larger than zero, this corresponds to irreversibility of a nonequilibrium system. By contrast, if there is no production entropy, $H_p = 0$, then this is a reversible, equilibrium system.

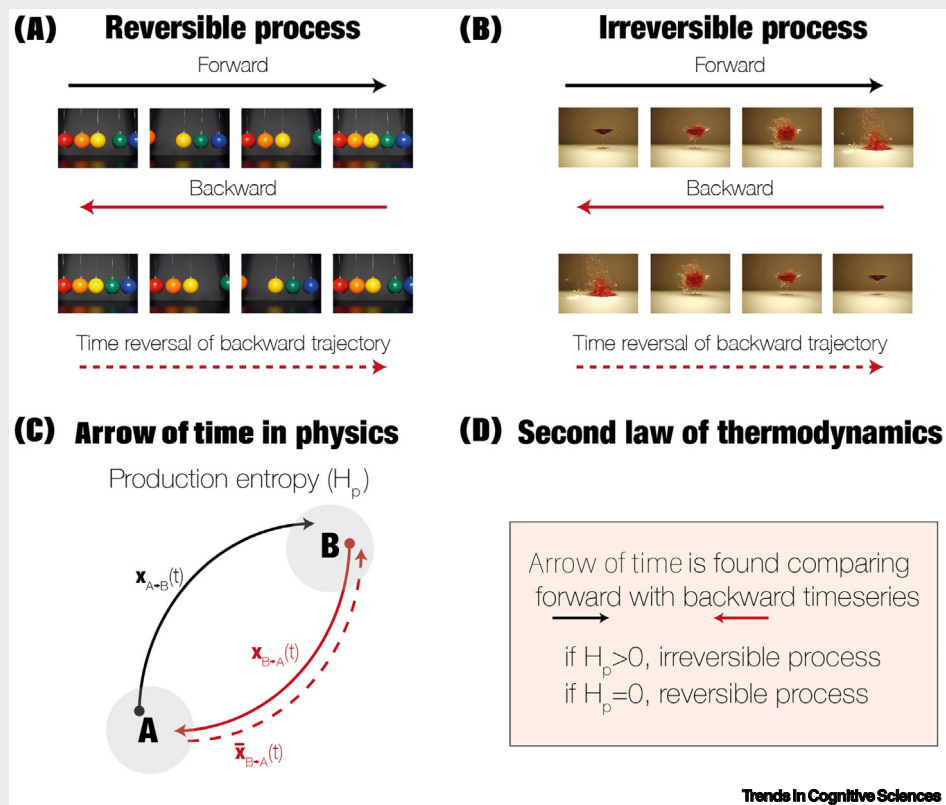


Figure 1. Physics and the arrow of time.

Unfortunately, as shown in a review by Hilgetag and Goulas [26], this rigour of analysis has not yet been applied to the brain. According to the authors, the term ‘hierarchy’ is currently not well defined in neuroscience. Bringing some order to the field, their careful analysis identifies four main

patterns used to characterise hierarchy: (i) topological projection sequences; (ii) spatially ordered changes (gradients) of features; (iii) progression of scales; and (iv) sorting of laminar projection patterns, of which they propose that the progression of scales appears the most fruitful way forward. This approach uses measures designed to capture the hierarchical, multiscale organisation of the brain by combining connectional sparsity at the global network level with network integration through connectivity that scales naturally from the local to the global level. In particular, the authors point out the usefulness of measures that capture the spacetime hierarchical scaling of the brain, that is, over space, with spatial encapsulation from ion channels and dendritic spines to neurons over local cortical column circuits of neurons to large-scale networks, and over time as temporal encapsulation of time scales and rhythms [27]. They also show that a useful measure of hierarchy is measuring segregation versus integration, which maximises the richness, for example, measured by the entropy, of potential functional interactions of local versus global access and control of networks [11,12].

Network theories are not very good at capturing the complexity of brain hierarchy since they all too often ignore the importance of time. However, they do support the general notion of a 'global workspace' orchestrating brain function, where information is integrated in a small group of brain regions before being broadcast to many other regions across the whole brain [28,29]. Global workspace can be thought of as a prototypical example of the orchestration of a hierarchical system, akin to how a conductor must control an orchestra, because, otherwise, the music almost always fails, as shown beautifully in Roberto Fellini's film *Prova d'orchestra*. Scientific studies have shown that this type of brain hierarchical organisation is efficient, robust, and largely fault tolerant [30–32].

Recently, transfer entropy, which is a generalised version of **Granger causality**, was used as a direct measure of hierarchy by estimating the information flow between brain signals. Transfer entropy was used on a large data set of over 1000 participants to identify a specific group of human brain regions constituting the global workspace used for the orchestration of the functional hierarchical organisation of the brain [33]. However, this transfer entropy framework is computationally expensive and requires long time series of fMRI data in large cohorts of participants, typically on the order of hundreds. Most data sets, including those from neuropsychiatric patients, tend to be much smaller. Therefore, a research focus has been to develop and use the more natural thermodynamics framework, which can not only provide a robust quantification of hierarchical organisation in smaller data sets in health and disease, but also allow for the discovery of the underlying **generative mechanisms**.

Hierarchy and thermodynamics

As shown in the preceding text, the hierarchy of a system can be determined by quantifying the asymmetry in the directionality of information flow. Conveniently, thermodynamics directly defines a measure of asymmetry for any physical system by the so-called 'breaking the detailed balance', which is a hallmark of nonequilibrium. This insight comes from the core idea of the second law of thermodynamics, namely, that a system will go from order to disorder over time [34,35]. Formally, thermodynamics relies on the concept of entropy, which is the level of disorder produced by irreversible (nonreversible) processes. However, the long history of thermodynamics has led to many different but related forms of entropy (Box 2).

Importantly, of these different forms of entropy, in terms of hierarchy, the concept of production entropy can directly quantify the level of irreversibility of a system and, therefore, the level of hierarchy. Yet, from a practical point of view, production entropy turns out to be hard to quantify in a high-dimensional system, such as the brain (although some progress has been made [23,36]).

Box 2. Entropy

Entropy is a rich concept with a long history, which has been immensely useful for quantifying physical systems in terms of the underlying disorder, uncertainty, transmission of information, and non-reversibility (arrow of time). All these concepts are different but related on a deep level. Briefly, entropy was first introduced in the 1850s by Rudolf Clausius as a measure of the disorder or randomness of a system [34]. Building on the seminal work of Sadi Carnot [35], he used the term to describe the amount of energy in a system that is unavailable to do work. As such, the second law of thermodynamics can elegantly be stated as an increase over time of the total entropy of a closed system.

This work gave birth to statistical physics during the 1870s, where Ludwig Boltzmann defined entropy as a measure of the number of possible microscopic configurations of a system that are consistent with a given macroscopic state: $S = k_B \ln \Omega$, where k_B denotes Boltzmann's constant and Ω is the number of microstates consistent with a given macroscopic equilibrium. This simple definition allowed him to show that the entropy of a system is related to the number of ways in which the particles in the system can be arranged and that the entropy of a system will increase as the number of possible configurations increases.

On a deep level, the measure of entropy as disorder is related to the randomness or uncertainty of information in a message, as shown in 1948 by Claude Shannon with his concept of information entropy [73] (Equation I):

$$H = -K \sum_{i=1}^k p(i) \ln p(i) \quad [I],$$

where K is a constant, determining the unit of measurement, stating that the information content of a message is related to the probability of each symbol in the message, $p(i)$, where the more uncertain or random the symbols are, the more information the message contains. This information-based definition of entropy directly leads to a definition of the mutual information (Equation II):

$$I(X; Y) = H(X) + H(Y) - H(X, Y) \quad [II].$$

As can be seen, this definition uses the information entropy to quantify the statistical dependence between two random variables, X and Y . This can then be used to define the transfer entropy from a variable X to Y as the conditional mutual information (Equation III):

$$T_{X \rightarrow Y} = I(Y_{i+1}; X^i | Y^i) = H(Y_{i+1} | Y^i) - H(Y_{i+1} | X^i, Y^i) \quad [III],$$

where Y_{i+1} is the value of Y at time point $i + 1$, and X^i indicates the past of X in a time window of length L up to, and including, the time point i , i.e., $X^i = [X_i X_{i-1} \dots X_{i-(L-1)}]$.

Adding to these definitions of entropy, central to thermodynamics is the production entropy, which is a measure of the reversibility and defined in Equation IV:

$$H_P = \sum_{i,j} P_{ij} \log \left(\frac{P_{ij}}{P_{ji}} \right) \quad [IV],$$

where P_{ij} is the probability of transition between states i at time t to j at time $t + 1$.

Overall, quantifying the production entropy in the brain is essentially detecting the arrow of time in brain signals, which is used in the Thermodynamics of Mind framework to quantify the hierarchical organisation of the underlying brain dynamics.

Instead, it is possible to estimate the irreversibility through another route, namely to compute the 'arrow of time' [37] to quantify the asymmetry of information flow in a complex system (Box 1).

Thus, quantifying hierarchy using thermodynamics is a matter of determining the 'arrow of time' or irreversibility in a system. The hierarchical organisation of a system is simply determined by the level of irreversibility. When a system is hierarchical, it is in nonequilibrium and, consequently, irreversible in time. Specifically, more hierarchical systems are more irreversible. At the other extreme, a system is flat and nonhierarchical when in equilibrium and reversible in time.

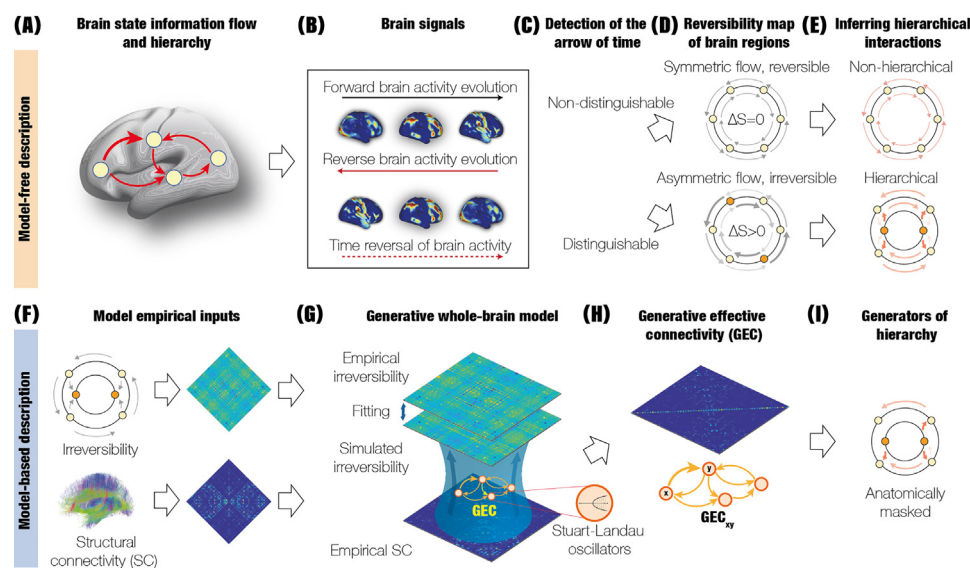
Overall, these insights gave rise to our novel Thermodynamics of Mind framework for advancing the understanding of brain dynamics. This framework has made it possible to reveal the underlying mechanisms of the hierarchical organisation of brain states, including wakefulness, sleep,

cognitive tasks (e.g., decision-making and working memory), drugs (anaesthesia and psychedelics), and disease (coma and neuropsychiatric disorders).

The Thermodynamics of Mind framework

Unlike other methods, the Thermodynamics of Mind framework can directly quantify the hierarchy of any brain state and provides insights into the mechanisms that generate this hierarchy. This has already shown important differences in cognition and brain states that other methods have failed to reveal.

Given the complexity of best characterising hierarchy through different measures of irreversibility, this means that there are multiple ways to implement the Thermodynamics of Mind framework. The most straightforward implementations directly quantify hierarchical changes from spatiotemporal whole-brain data in a 'model-free' way. Figure 1A–E shows how the thermodynamic approach can quantify hierarchy through the level of irreversibility for a given brain state. Crucially, this process relies on distinguishing the differences in the arrow of time for forward and backward time series of brain signals [38–40].



Trends In Cognitive Sciences

Figure 1. Thermodynamics of Mind framework. The figure describes how the principles of thermodynamics can reveal the hierarchy of brain states by quantifying the asymmetry in the information flow. (A) Hierarchical organisation can be found from the asymmetry in the directionality of information flow, called 'breaking of the detailed balance' in thermodynamics. (B) Discovering the directionality, or the 'arrow of time', requires forward time series of brain signals, as well as the backward time series (by artificially reverting the time ordering). (C) The detection of irreversibility requires distinguishing between these forward and backward time series. (D) The irreversibility of a system is spread between two extreme cases: fully reversible and symmetric between all nodes (top) or irreversible, with asymmetric flow between nodes breaking the detailed balance (bottom). (E) In terms of hierarchy, these two extremes translate into different hierarchies: uniform symmetric flow leads to a flat hierarchy (top), while asymmetric flow leads to a hierarchical system (top), with the two orange regions (in the inner circle) at the top of the hierarchy. Thus, the arrow of time provides a convenient measure of both the flow and hierarchy of any dynamical system. (F) Beyond this, whole-brain modelling can be used to identify the causal, mechanistic generators of the functional hierarchy. The model integrates the anatomical structural connectivity and functional dynamics quantified using the model-free irreversibility estimates. (G) The whole-brain model uses local dynamics, for example, Stuart–Landau oscillators, to fit the empirical irreversibility. (H) The optimisation leads to the best possible generative effective connectivity (GEC). (I) In turn, the GEC provides an estimate of the generators giving rise to the generative hierarchy of a given brain state, which can be characterised using trophic coherence capturing cycle structure, stability, and percolation [70].

It is also possible to move beyond the model-free approach to reveal the underlying generative mechanisms of hierarchy (Figure 1F–I). This is achieved by constructing a generative **whole-brain model** integrating anatomical structural connectivity and functional dynamics [41]. Fitting the model to the irreversibility creates the so-called ‘generative effective connectivity’, which is a matrix revealing the causal, mechanistic generators of the specific functional hierarchy.

Orchestration of information flow in brain states and cognition

Traditional neuroimaging methods have revealed some of the brain dynamics involved in human cognition. Yet, what has been missing is a deeper understanding of how complex cognition is orchestrated through different patterns of information flow between brain networks. Quantifying these subtle differences has proven a challenge for traditional analysis methods. However, the Thermodynamics of Mind framework has revealed significant differences in orchestration for even quite similar cognitive tasks.

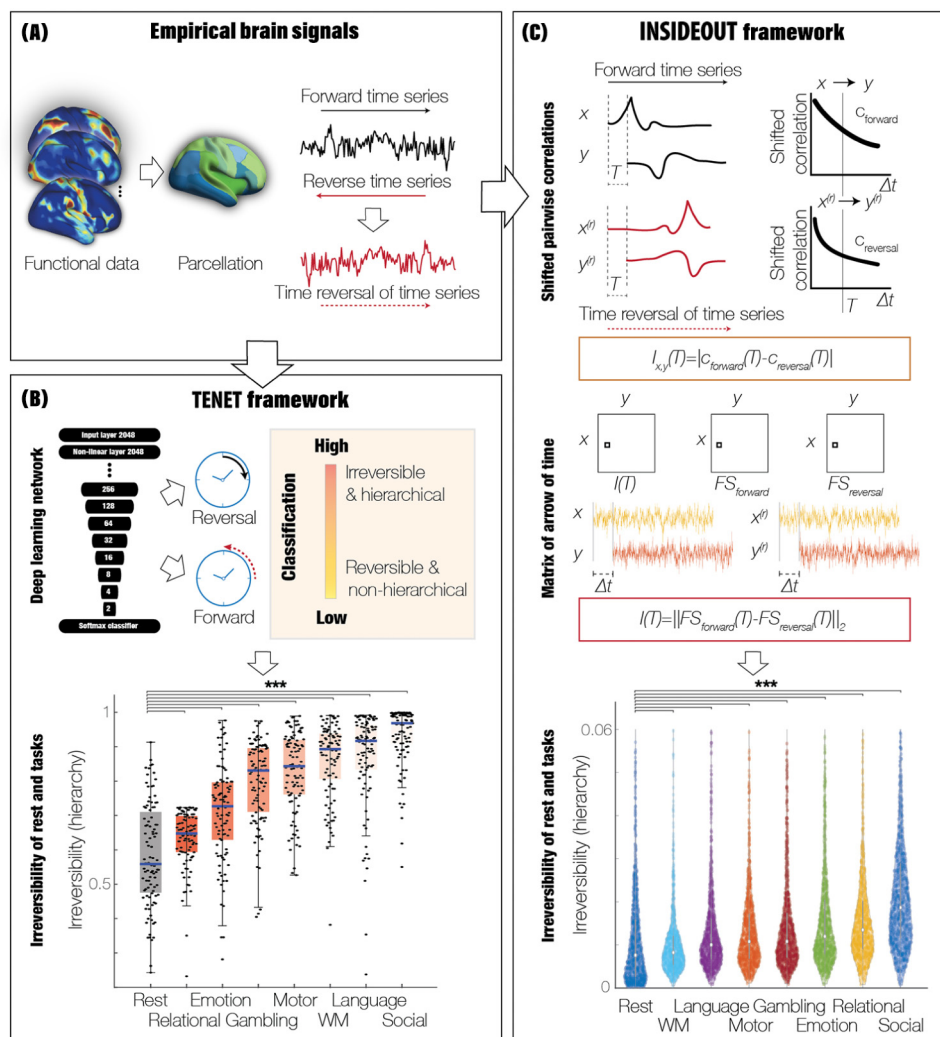
For example, one recent study used a machine learning (Temporal Evolution NETWORK, TENET) implementation of the quantification of irreversibility for fMRI data and found that resting state has significantly lower irreversibility and flatter hierarchy than when performing seven cognitive tasks (Figure 2A–C) [38]. The higher level of irreversibility observed during tasks reflects the increase in the hierarchical organisation needed for the specific computations, reflecting the increase in asymmetrical information flow in tasks compared with resting state. TENET was also applied to large-scale fMRI data from neuropsychiatric patients with attention deficit-hyperactivity disorder (ADHD), bipolar disorder, and schizophrenia, which revealed overall lower irreversibility during resting compared with in control participants [38], suggestive of significant problems with orchestration. This was further elucidated through the significant local heterogeneous node-level changes in the different disorders.

Still, while the TENET implementation is powerful and robust in terms of the underlying machine learning framework [42], it is also computationally expensive and requires large data sets. To mitigate these problems, the INSIDEOUT implementation was developed, which is a robust way of capturing the ‘inside out’ balance of intrinsic (INSIDE) and extrinsic (OUT) brain dynamics. This is achieved by directly estimating the arrow of time in brain signals (Figure 2D) [39]. Briefly, the main idea of INSIDEOUT is to use the simplicity of time-shifted correlation matrices to provide a quantification of irreversibility and, consequently, the degree of nonequilibrium in the brain dynamics of different brain states.

When applied to the same fMRI data, INSIDEOUT gives very similar results to TENET when measuring the irreversibility of cognitive states compared with resting state (Figure 2E) [39]. Crucially, however, unlike TENET, INSIDEOUT allows for the estimation of precise signatures in much smaller data sets, such as electrocorticography brain data from, for example, individual non-human primates [43,44]. When comparing three radically different brain states of awake, deep sleep, and anaesthesia in non-human primates, this revealed a significantly different hierarchical organisation in each brain state in terms of non-reversibility and hierarchy. Potentially, this may be a signature of conscious awareness, demonstrating a flattening of the hierarchical organisation as the level of consciousness decreases [39].

Watching movies: perched between rest and cognition

Still, truly understanding hierarchical brain organisation requires moving beyond simply quantifying irreversibility in brain states to building whole-brain models that can explain the underlying data. Such mechanistic models are both informative of the information flow and hierarchical



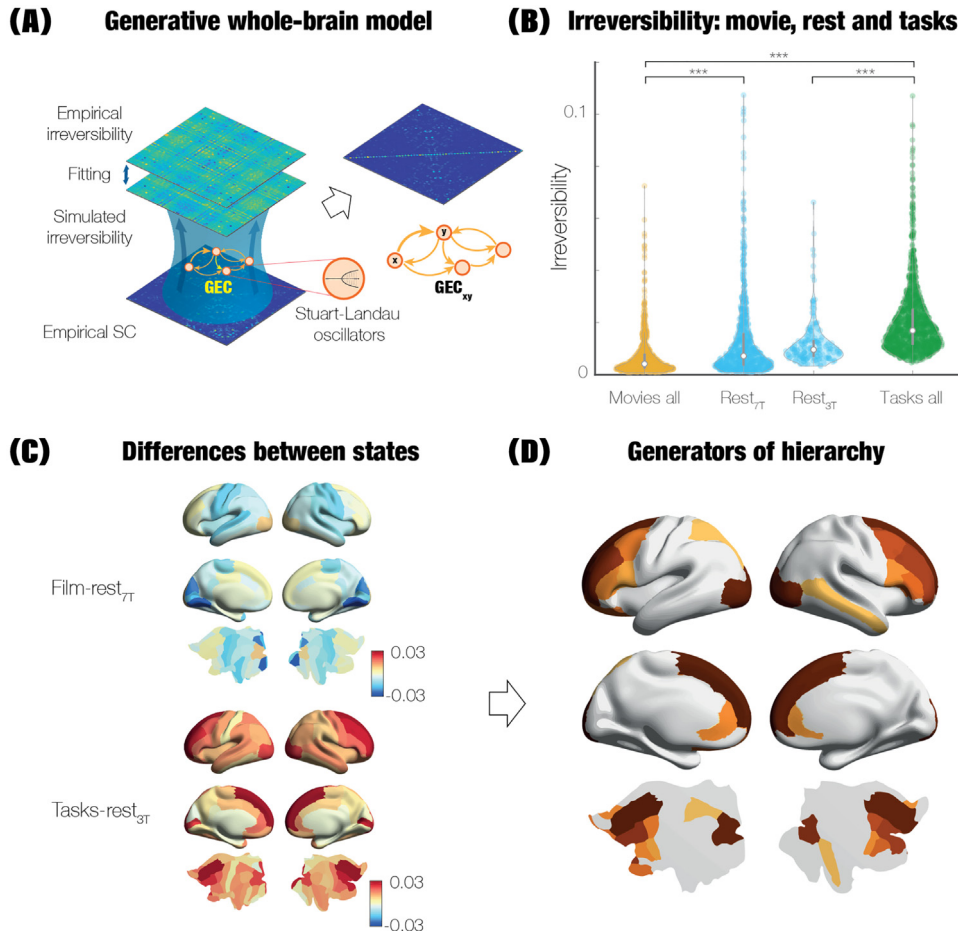
Trends In Cognitive Sciences

Figure 2. Hierarchy in brain dynamics. The figure shows two implementations of the Thermodynamics of Mind framework, which can establish the arrow of time in the brain and reveal the brain hierarchy. (A) Brain signals are extracted in a brain parcellation, providing forward and backward time series (by artificially reverting the time ordering), which are then used to establish irreversibility. (B) Temporal Evolution deep learning NETWORK (TENET) uses deep learning on these time series for training and, when subsequently testing new data, the level of classification performance provides the level of irreversibility [38] (top panel). Using TENET on large-scale neuroimaging data from over 1000 participants in the Human Connectome Project (HCP) showed that the resting state had a lower non-reversibility and flatter hierarchy than when participants performed seven cognitive tasks (bottom panel). (C) Similarly, the INSIDEOUT framework also uses forward and backward time series to establish the arrow of time information flow. The framework generates the irreversibility by comparing shifted pairwise correlations between forward and backward time series [39,40]. These pairwise measures are then combined into a full matrix from which the hierarchy can be established. Applying INSIDEOUT implementation to the large HCP data set revealed similar results to TENET, namely that resting state has a lower non-reversibility and flatter hierarchy compared with when seven tasks were performed. Abbreviation: WM, working memory.

organisation in brain states and can capture the mechanisms generating a given brain hierarchy, which is needed to resolve the causal attribution [15].

In fact, a whole-brain model implementation of the Thermodynamics of Mind framework was able to shed light on a fundamental question in neuroscience, namely why our subjective

experience of watching the naturalistic, multimodal dynamics of film is different and much more pleasant than our everyday resting experiences of both mind wandering and solving difficult tasks [41]. This is important because naturalistic films have been proposed as a better alternative to resting state for investigating younger and clinical populations, especially given that using film results in higher test–retest reliability [45].



Trends in Cognitive Sciences

Figure 3. Thermodynamic insights into movie watching. Whole-brain modelling implementation of the Thermodynamics of Mind theory has provided insights into why our subjective experience of watching the naturalistic, multimodal dynamics of film is so highly motivating, soothing, and entirely different from our usual everyday resting experience of mind wandering [41]. (A) The implementation built a generative whole-brain model fitting the model-free estimates of irreversibility in a large group of participants watching naturalistic movies and resting of performing seven tasks in a 7 Tesla MRI scanner [41]. (B) Again, the results showed highest irreversibility in tasks and significantly lower irreversibility in rest. However, interestingly naturalistic movies resulted in significantly lower irreversibility than either tasks or at rest and, thus, a flatter brain hierarchy (significant differences indicated by ***). The whole-brain model was able to capture the main mechanistic drivers of these changes in hierarchy through generative effective connectivity (GEC), which provides the underlying causal mechanisms for a given state. (C) The causal brain generators of movie watching are shown in renderings of the differences of the sum of the incoming receivers and outgoing drivers of the GEC matrices for (1) naturalistic film versus rest and (2) rest versus tasks. (D) Most importantly, rendering the intersection of the top 50% regions of these two contrasts shows that the prefrontal cortex is the primary driver for orchestrating computation in the brain (but with some parietal, visual, and temporal regions). Overall, this provides important quantifications of the causal mechanisms underlying complex changes in brain hierarchy.

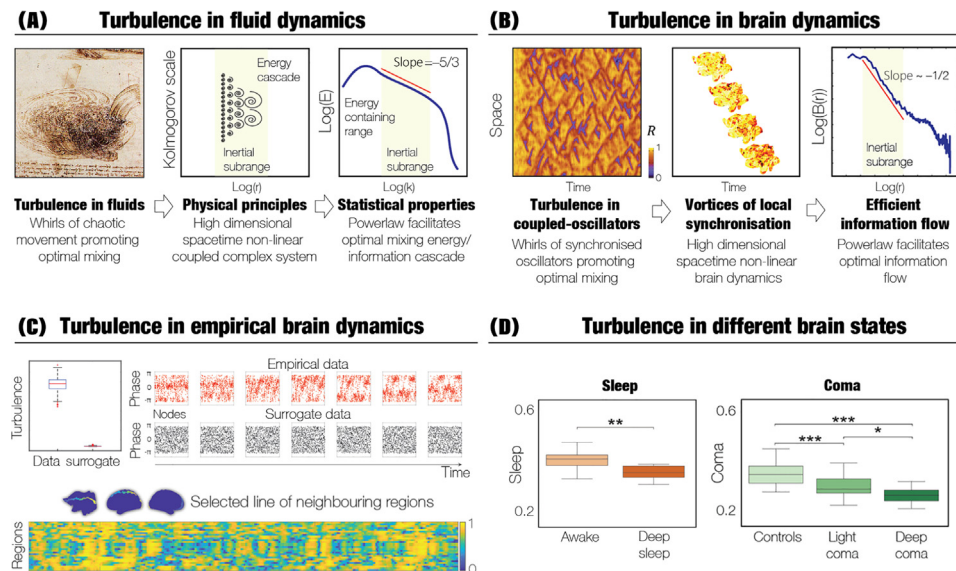
Figure 3B shows the result of watching naturalistic movies, which is associated with significantly lower levels of irreversibility compared with both resting and tasks measured with fMRI [41]. The flatter hierarchy during movie watching could perhaps be linked to why watching films is such a favourite relaxing pastime worldwide. Specifically, naturalistic films appear to offer a moment's respite from the thermodynamic rat race of survival. Movie watching provides a desirable audio-visual narrative where the necessary computation is minimal, which is very different from both solving demanding tasks and simply resting. Interestingly, resting is not particularly restful for most people, with authors such as Killingsworth and Gilbert showing that the introspection and mind-wandering state rarely lead to a happy mind [46].

Fast brain processing despite slowness

Beyond providing a deeper understanding of information flow in the human brain, the Thermodynamics of Mind framework also helps solve another key problem in neuroscience, namely how the brain can survive given the relative slowness of the underlying signals [47]. Indeed, how the brain overcomes the limitations of speed for information transfer across spacetime has long been a conundrum in neuroscience.

Thermodynamics can also provide an answer to this unsolved problem. Recent research showed that the answer lies in turbulence (Figure 4), the irreversible dynamical regime taking place far from equilibrium while showing strong time asymmetry. Originally coined as 'turbolenza' by Leonardo da Vinci over half a millennium ago [48] and subsequently developed by many mathematicians, turbulence is ubiquitous in nature as an essential dynamical regime facilitating efficient energy and information transfer across spatiotemporal scales [49]. Importantly, Andrey Kolmogorov demonstrated the efficiency of turbulence by finding a spatial power scaling law in fluid dynamics [50,51] (Figure 4A). Beyond the limited domain of fluid dynamics, turbulence is also found in other physical systems, including coupled oscillators [52] and brains [53,54] (Figure 4B,C). Modelling the brain as coupled oscillators, the level of turbulence has been shown to vary between different brain states, including different forms of coma, sleep [55] (Figure 4D), and psychedelics [56]. Remarkably, brain dynamics has also been found to exhibit a similar turbulent **power law**, strongly suggesting the presence of a cascade of efficient information processing across scales [53] and, more recently, evidence has been found of higher-order structure functions demonstrating turbulence [57]. Interestingly, turbulence is spatiotemporal chaos and can produce stochastic features similar to what some may classify as 'noise' but what are in fact bound inextricably to the brain signals [58].

Important to the question of how the brain overcomes the slowness of the local signals, turbulence has been demonstrated not only in whole-brain networks derived from slow haemodynamic signals measured with fMRI, but also in fast local hippocampal circuits [59]. Recently, a fast whole-brain study using magnetoencephalography provided insight into how turbulence could provide the skeleton underlying efficient spatiotemporal information transfer required for survival [47]. This study showed that, similar to the insights of Kolmogorov [50,51], the cascading whirls of turbulence are at the root of fast, efficient information transfer. Indeed, the presence of a spatial power-scaling law is a hallmark of turbulence and provides a mathematical description of the concept of cascaded eddies [60]. This conforms remarkably with Leonardo's observation that the constriction of circumference toward the centre of the vortex is more rapid than the diminution of the impetus of the water, which is why the water revolves faster near the centre. As such, turbulence promises to be a highly sensitive thermodynamic principle by which the brain is able to compute the time-critical behaviour needed for survival.



Trends in Cognitive Sciences

Figure 4. Turbulence provides efficient, fast information transfer in the brain. The Thermodynamics of Mind framework can also describe turbulence, which is a dynamical regime taking place far from equilibrium. (A) Leonardo da Vinci coined the term ‘turbolenza’ for the seemingly chaotic dynamics of fluids. The physical principles giving rise to turbulence are given by high-dimensional spacetime nonlinear coupled systems [48]. The excellent mixing capabilities of fluid turbulence come from the energy cascade turning large whirls into smaller whirls and eventually energy dissipation. Furthermore, the turbulent energy cascade has been shown to be highly efficient across scales, as evidenced by a power law. (B) Moving beyond fluid dynamics, empirical brain dynamics in resting state data from 1000 healthy participants were recently shown to exhibit turbulence [53]. This shows highly variable, local synchronisation vortices across time and space. Equally, the turbulent brain regime also gives rise to an efficient information cascade obeying a power law [53]. Furthermore, Hopf whole-brain models can be used to gain a causal understanding of the underlying mechanisms by fitting both turbulence and the empirical neuroimaging data at the same working point [54]. (C) Turbulence is found in resting state but not in carefully matched surrogate data, which can be seen in consecutive snapshots over time of the phases of all brain regions for the empirical data (red, top) and the surrogate data (black, bottom). The bottom panel visualises the turbulent spatiotemporal evolution of neighbourhood-dependent synchronisation in a time-evolving plot of 26 neighbouring parcels. (D) Turbulence robustly distinguishes radically different brain states, such as deep sleep and coma, which show significantly lower levels of turbulence compared with resting state, reflecting the reduced level of information flow in these states [55].

Future avenues for research

The Thermodynamics of Mind framework provides a key stepping stone toward a fundamental theory for describing the orchestration of hierarchical brain function. It takes inspiration from the long illustrious history of thermodynamics, which includes important contributions by Nobel Prize-winning physicists Albert Einstein and Erwin Schrödinger, perhaps even as important as their seminal work in quantum mechanics and relativity.

Schrödinger and Einstein were close friends but quarrelled over the interpretation of quantum mechanics, especially in their 1935 discussions about Schrödinger’s cat, which, paradoxically, may be considered simultaneously both alive and dead. Yet, both agreed on the pre-eminence of thermodynamics as a fundamental and fruitful framework for understanding the flow and mechanisms of any hierarchical physical system, including biological living organisms. In fact, Einstein’s Nobel Prize-winning work in his *annus mirabilis* of 1905 was based firmly on thermodynamics. In his later exile in Ireland, Schrödinger went even further and wrote his important book *What Is Life?*, in which he proposed that thermodynamics, specifically the arrow of time, or non-reversibility, are crucial elements for sustaining life. In particular, he wrote ‘How does the living organism avoid decay? ... By eating, drinking, breathing and ... assimilating. The technical term is metabolism’ [61].

Yet, neither Schrödinger nor Einstein thought to extend the framework of thermodynamics to the brain. This is perhaps not surprising given that the scientific study of the brain was still in its infancy at that time. However, neuroscience has now produced an abundance of empirical data calling out for more fundamental theories of brain function. In some ways, this parallels the abundance of empirical data from physical systems that inspired both Schrödinger and Einstein at the beginning of the 20th century. The Thermodynamics of Mind framework shows how life sciences can benefit from the important insights from physicists.

In our quest for a general theory, we continue to be inspired by the success of modern physics and its ability to model the constitutive elements of a system and to exhaustively perturb these to discover the emergence of the underlying dynamics [62]. Turbulence, whether in oscillators or fluids, is an excellent example of how a detailed understanding has come about through careful modelling of the statistical properties of the necessary and sufficient elements over multiple scales [63]. Other future fertile avenues of research include using the thermodynamic concept of balancing friction and thermal noise, that is, the balancing forces of dissipation and spontaneous fluctuations, which are at the heart of Einstein's theory of Brownian motion [64]. This is commonly referred to as the **fluctuation–dissipation theorem (FDT)**, which has been highly successful in describing many different kinds of physical system in both equilibrium and nonequilibrium [65].

In terms of using similar perturbative principles to better understand the fundamental brain principles, one future avenue could be to implement a model of FDT to directly combine model-free and model-based approaches (see [Outstanding questions](#)). This model could use a perturbation to reveal the state of nonequilibrium of the brain. As such, this would be another important instantiation of the Thermodynamics of Mind framework unifying model-free and model-based approaches. This could be more sensitive for describing the hierarchy in a brain state. More specifically, a generative perturbative whole-brain model should be able to estimate the violation of FDT in empirical neuroimaging data from humans in different brain states (such as in wakefulness, cognitive tasks, and deep sleep). This perturbative model-based approach would go beyond any model-free analysis of unperturbed brain states. In fact, this approach could use Onsager's regression principle leading to a simple derivation of FDT [66–68]. This derivation holds that, when a system in an initial equilibrium state is driven by an external perturbation to a final equilibrium state, then the evolution of the system from an initial to a final state can be treated as a spontaneous equilibrium fluctuation. Crucially, beyond this equilibrium state, the FDT framework would allow for a characterisation of the level of nonequilibrium by simply computing the violation of the FDT [69]. Future research should systematically investigate this important question.

Looking further ahead in refining and expanding our Thermodynamics of Mind framework, one avenue could be to make thermodynamics-based whole-brain models of non-human animals under anaesthesia and predict how to intervene to awaken them. More generally, we are also looking for better descriptions of the time-evolving temporal hierarchy, since this would provide a way to get at the fundamental computations in the brain. Methods from thermodynamics promise to provide exactly such precise descriptions of the changing of hierarchy over time at the whole-brain level by estimating the evolving low-dimensional manifold of brain activity; potentially providing a novel perspective on computation and learning that could significantly add to the rich literature on cognition.

Concluding remarks

Overall, the general Thermodynamics of Mind framework holds great promise for revealing the underlying principles of hierarchy in shaping the nonequilibrium of the human brain. Specifically, various implementations of the framework have already shed new light on the changes in

Outstanding questions

What other thermodynamic concepts may help quantify brain hierarchy? One candidate is FDT, which describes balancing forces of dissipation and spontaneous fluctuations. Using FDT with a generative, perturbative whole-brain model can estimate the violation of FDT in different brain states, such as in wakefulness, cognitive tasks, and deep sleep, and bring new insights into the causal orchestration of hierarchy in these states.

How can the Thermodynamics of Mind framework best be refined and expanded to provide better descriptions of time-evolving temporal hierarchy? This could provide a principled way to describe the fundamental brain computations. In particular, this could provide a quantification of the changing of hierarchy over time at the whole-brain level by estimating the evolving low-dimensional manifold of brain activity and provide a novel perspective on computation and learning.

What are the underlying principles of hierarchy in shaping the nonequilibrium of the human brain? Implementations of the framework have already shed new light on the changes in orchestration and hierarchical organisation in health and may in the future help to better understand the breakdown in neuropsychiatric disease.

What novel predictions can the Thermodynamics of Mind framework make? The framework predicts that the treatment of a patient with depression with different pharmacological treatments will lead to different hierarchical reconfigurations. This could explain why some pharmacological treatments are more effective and have more severe side effects than do others. The framework will be able to distinguish treatment and response interaction effects when comparing, for example, the treatment of depression with standard SSRIs and psychedelics.

orchestration and hierarchical organisation in health and may help understand the breakdown in neuropsychiatric disease. While the Thermodynamics of Mind is not a general theory of brain function but merely a step toward this, it yields clear predictive power. For example, the framework predicts that, when the brain state of a patient with depression is being changed by different pharmacological treatments, this could lead to different paths of hierarchical reconfigurations. In other words, the framework predicts that patients who get better from a treatment could have different hierarchical reorganisations. This could also explain why some pharmacological treatments have more severe side effects than do others. A pertinent example would be to compare the treatment of depression with psychedelics and selective serotonin-reuptake inhibitors (SSRIs) and to distinguish treatment and response interaction effects. Still, beyond such practical predictions, Thermodynamics of Mind may hold the key to eventually discovering the very nature of complex distributed brain computations and, in a deep sense, provide an answer to Schrödinger's deceptively simple question, 'what is life?' [61].

Acknowledgments

We would like to thank Shamil Chandaria for his very insightful suggestions on an earlier version of this manuscript. G.D. is supported by Grant PID2022-136216NB-I00 funded by MICIU/AEI/10.13039/501100011033 and by 'ERDF A way of making Europe', ERDF, EU, Project NEurological MEchanisms of Injury, and Sleep-like cellular dynamics (NEMESIS) (101071900) funded by the EU ERC Synergy Horizon Europe, and AGAUR research support grant (2021 SGR 00917) funded by the Department of Research and Universities of the Generalitat of Catalunya. Y.S.P. is supported by the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant 896354 and NEMESIS (101071900), funded by the EU ERC Synergy Horizon Europe. M.L.K. is supported by the Centre for Eudaimonia and Human Flourishing (funded by the Pettit and Carlsberg Foundations) and Center for Music in the Brain (funded by the Danish National Research Foundation, DNRF117).

Declaration of interests

None declared by authors.

References

1. Cotterill, R.M.J. (2002) *Biophysics. An introduction*, Wiley
2. Itoh, K. *et al.* (2022) Cerebral cortical processing time is elongated in human brain evolution. *Sci. Rep.* 12, 1103
3. Hodgkin, A.L. and Huxley, A.F. (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol. (London)* 117, 500–544
4. Mesulam, M.M. (1998) From sensation to cognition. *Brain* 121, 1013–1052
5. Felleman, D.J. and Van Essen, D.C. (1991) Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* 1, 1–47
6. Zeki, S. and Shipp, S. (1988) The functional logic of cortical connections. *Nature* 335, 311–317
7. Kringelbach, M.L. *et al.* (2020) Dynamic coupling of whole-brain neuronal and neurotransmitter systems. *Proc. Natl. Acad. Sci. U. S. A.* 117, 9566–9576
8. Deco, G. *et al.* (2018) Whole-brain multimodal neuroimaging model using serotonin receptor maps explains non-linear functional effects of LSD. *Curr. Biol.* 28, 3065–3074
9. Shine, J.M. *et al.* (2021) Computational models link cellular mechanisms of neuromodulation to large-scale neural dynamics. *Nat. Neurosci.* 24, 765–776
10. Sporns, O. *et al.* (2000) Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices. *Cereb. Cortex* 10, 127–141
11. Tononi, G. *et al.* (1994) A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proc. Natl. Acad. Sci. U. S. A.* 91, 5033–5037
12. Deco, G. *et al.* (2015) Rethinking segregation and integration: contributions of whole-brain modelling. *Nat. Rev. Neurosci.* 16, 430–439
13. Gervasoni, D. *et al.* (2004) Global forebrain dynamics predict rat behavioral states and their transitions. *J. Neurosci.* 24, 11137–11147
14. Northoff, G. (2013) What the brain's intrinsic activity can tell us about consciousness? A tri-dimensional view. *Neurosci. Biobehav. Rev.* 37, 726–738
15. Kringelbach, M.L. and Deco, G. (2020) Brain states and transitions: insights from computational neuroscience. *Cell Rep.* 32, 108128
16. Rao, R.P. and Ballard, D.H. (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87
17. Golesorkhi, M. *et al.* (2021) The brain and its time: intrinsic neural timescales are key for input processing. *Commun. Biol.* 4, 970
18. Wolff, A. *et al.* (2022) Intrinsic neural timescales: temporal integration and segregation. *Trends Cogn. Sci.* 26, 159–173
19. Golesorkhi, M. *et al.* (2022) From temporal to spatial topography: hierarchy of neural dynamics in higher- and lower-order networks shapes their complexity. *Cereb. Cortex* 32, 5637–5653
20. Golesorkhi, M. *et al.* (2021) Temporal hierarchy of intrinsic neural timescales converges with spatial core-periphery organization. *Commun. Biol.* 4, 277
21. Knill, D.C. and Pouget, A. (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27, 712–719
22. Friston, K. (2010) The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11, 127–138
23. Lynn, C.W. *et al.* (2021) Broken detailed balance and entropy production in the human brain. *Proc. Natl. Acad. Sci. U. S. A.* 118, e2109889118
24. Djeraba, C. and Simovici, D.A. (2008) *Mathematical Tools for Data Mining: Set Theory, Partial Orders, Combinatorics*, Springer
25. Evans, F.C. (1956) Ecosystem as the basic unit in ecology. *Science* 123, 1127–1128
26. Hilgetag, C.C. and Goulas, A. (2020) 'Hierarchy' in the organization of brain networks. *Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci.* 375, 20190319

27. Buzsaki, G. (2006) *Rhythms of the Brain*, Oxford University Press
28. Baars, B.J. (1989) *A Cognitive Theory of Consciousness*, Cambridge University Press
29. Dehaene, S. et al. (1998) A neuronal model of a global workspace in effortful cognitive tasks. *Proc. Natl. Acad. Sci. U. S. A.* 95, 14529–14534
30. Bullmore, E. and Sporns, O. (2012) The economy of brain network organization. *Nat. Rev. Neurosci.* 13, 336–349
31. Honey, C.J. and Sporns, O. (2008) Dynamical consequences of lesions in cortical networks. *Hum. Brain Mapp.* 29, 802–809
32. Alstott, J. et al. (2009) Modeling the impact of lesions in the human brain. *PLoS Comput. Biol.* 5, e1000408
33. Deco, G. et al. (2021) Revisiting the Global Workspace orchestrating the hierarchical organisation of the human brain. *Nat. Hum. Behav.* 5, 497–511
34. Clausius, R. (1865) Ueber verschiedene für die Anwendung bequeme Formen der Hauptgleichungen der mechanischen Wärmetheorie (Vorgetragen in der naturforsch. Gesellschaft zu Zürich den 24. April 1865). *Ann. Phys. Chem.* 125, 353–400
35. Carnot, S. (1824) *Reflections on the Motive Power of Fire and on Machines Fitted to Develop that Power*, Bachelier, Paris, p. 108
36. Sanz Perl, Y. et al. (2021) Non-equilibrium brain dynamics as a signature of consciousness. *Phys. Rev. E* 104, 014411
37. Eddington, A.S. (1928) *The Nature of the Physical World*, Macmillan
38. Deco, G. et al. (2023) The arrow of time of brain signals in cognition: potential intriguing role of parts of the default mode network. *Netw. Neurosci.* 7, 966–998
39. Deco, G. et al. (2022) The INSIDEOUT framework provides precise signatures of the balance of intrinsic and extrinsic dynamics in brain states. *Commun. Biol.* 5, 572
40. G-Guzman, E. et al. (2023) The lack of temporal brain dynamics asymmetry as a signature of impaired consciousness states. *Interface Focus* 13, 20220086
41. Kringelbach, M.L. et al. (2023) Toward naturalistic neuroscience: mechanisms underlying the flattening of brain hierarchy in movie-watching compared to rest and task. *Sci. Adv.* 9, eade6049
42. de la Fuente, L. et al. (2022) Temporal irreversibility of neural dynamics as a signature of consciousness. *Cereb. Cortex* 33, 1856–1865
43. Nagasaka, Y. et al. (2011) Multidimensional recording (MDR) and data sharing: an ecological open research and educational platform for neuroscience. *PLoS ONE* 6, e22561
44. Yanagawa, T. et al. (2013) Large-scale information flow in conscious and unconscious states: an ECoG study in monkeys. *PLoS ONE* 8, e80845
45. Vanderwal, T. et al. (2019) Movies in the magnet: naturalistic paradigms in developmental functional neuroimaging. *Dev. Cogn. Neurosci.* 36, 100600
46. Killingsworth, M.A. and Gilbert, D.T. (2010) A wandering mind is an unhappy mind. *Science* 330, 932
47. Deco, G. et al. (2023) The effect of turbulence in brain dynamics information transfer measured with magnetoencephalography. *Commun. Phys.* 6, 74
48. Deco, G. et al. (2021) Leonardo da Vinci and the search for order in neuroscience. *Curr. Biol.* 31, R704–ER709
49. Cross, M.C. and Hohenberg, P.C. (1993) Pattern formation outside of equilibrium. *Rev. Mod. Phys.* 65, 851–1112
50. Kolmogorov, A.N. (1941) The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *Proc. USSR Acad. Sci. (Atmos. Ocean. Phys.)* 30, 299–303
51. Kolmogorov, A.N. (1941) Dissipation of energy in locally isotropic turbulence. *Proc. USSR Acad. Sci. (in Russian)* 32, 16–18
52. Kuramoto, Y. (1984) *Chemical Oscillations, Waves, and Turbulence*, Springer-Verlag
53. Deco, G. and Kringelbach, M.L. (2020) Turbulent-like dynamics in the human brain. *Cell Rep.* 33, 108471
54. Deco, G. et al. (2021) Rare long-range cortical connections enhance human information processing. *Curr. Biol.* 31, 1–13
55. Eschrichs, A. et al. (2022) Unifying turbulent dynamics framework distinguishes different brain states. *Commun. Biol.* 5, 638
56. Cruzat, J. et al. (2022) Effects of classic psychedelic drugs on turbulent signatures in brain dynamics. *Netw. Neurosci.* 6, 1104–1124
57. Perl, Y.S. et al. (2023) Scaling of whole-brain dynamics reproduced by high-order moments of turbulence indicators. *Phys. Rev. Res.* 5, 033183
58. Uddin, L.Q. (2020) Bring the noise: reconceptualizing spontaneous neural activity. *Trends Cogn. Sci.* 24, 734–746
59. Sheremet, A. et al. (2019) Wave turbulence and energy cascade in the hippocampus. *Front. Syst. Neurosci.* 12, 62
60. Richardson, L.F. (1922) *Weather Prediction by Numerical Process*, Cambridge University Press
61. Schrödinger, E. (1944) *What Is Life? The Physical Aspect of the Living Cell*, Cambridge University Press
62. Feynman, R.P. et al. (2005) *The Feynman Lectures on Physics Including Feynman's Tips on Physics: The Definitive and Extended Edition* (2nd edn), Addison-Wesley
63. Frisch, U. (1995) *Turbulence: The Legacy of A. N. Kolmogorov*, Cambridge University Press
64. Einstein, A. (1905) Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen [On the Movement of Small Particles Suspended in Stationary Liquids Required by the Molecular-Kinetic Theory of Heat]. *Ann. Phys.* 322, 549–560
65. Marconi, U.M.B. et al. (2008) Fluctuation–dissipation: response theory in statistical physics. *Phys. Rep.* 461, 111–195
66. Crisanti, A. and Ritort, F. (2003) Violation of the fluctuation–dissipation theorem in glassy systems: basic notions and the numerical evidence. *J. Phys. A Math. Gen.* 36, R181
67. Onsager, L. (1931) Reciprocal relations in irreversible processes. II. *Phys. Rev.* 38, 2265–2279
68. Onsager, L. (1931) Reciprocal relations in irreversible processes. I. *Phys. Rev.* 37, 405–426
69. Deco, G. et al. (2023) Violations of the fluctuation-dissipation theorem reveal distinct non-equilibrium dynamics of brain states. *Phys. Rev. E* 108, 064410
70. MacKay, R.S. et al. (2020) How directed is a directed network? *R. Soc. Open Sci.* 7, 201138
71. Jarzynski, C. (2011) Equalities and inequalities: irreversibility and the second law of thermodynamics at the nanoscale. *Annu. Rev. Condens. Matter Phys.* 2, 329–351
72. Seif, A. et al. (2021) Machine learning the thermodynamic arrow of time. *Nat. Phys.* 17, 105–113
73. Shannon, C.E. (1948) A mathematical theory of communication. *Bell Syst. Tech. J.* 27, 379–423 623–656