

TOPICAL REVIEW • OPEN ACCESS

The free energy principle induces neuromorphic development

To cite this article: Chris Fields *et al* 2022 *Neuromorph. Comput. Eng.* **2** 042002

View the [article online](#) for updates and enhancements.

You may also like

- [Fluctuation-driven initialization for spiking neural network training](#)
Julian Rossbroich, Julia Gygax and Friedemann Zenke
- [Edge to quantum: hybrid quantum-spiking neural network image classifier](#)
A Ajayan and A P James
- [A mixed-signal oscillatory neural network for scalable analog computations in phase domain](#)
Corentin Delacour, Stefania Carapezzi, Gabriele Boschetto et al.



TOPICAL REVIEW

OPEN ACCESS

RECEIVED
20 July 2022REVISED
2 November 2022ACCEPTED FOR PUBLICATION
1 December 2022PUBLISHED
16 December 2022

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



The free energy principle induces neuromorphic development

Chris Fields^{1,2,*} , Karl Friston³ , James F Glazebrook^{4,5} , Michael Levin^{2,6}
and Antonino Marciano^{7,8,9} ¹ 23 Rue des Lavandières, 11160 Caunes Minervois, France² Allen Discovery Center at Tufts University, Medford, MA 02155, United States of America³ Wellcome Centre for Human Neuroimaging, University College London, London WC1N 3AR, United Kingdom⁴ Department of Mathematics and Computer Science, Eastern Illinois University, Charleston, IL 61920 United States of America⁵ Adjunct Faculty, Department of Mathematics, University of Illinois at Urbana-Champaign, Urbana, IL 61801 United States of America⁶ Wyss Institute for Biologically Inspired Engineering at Harvard University, Boston, MA 02115, United States of America⁷ Center for Field Theory and Particle Physics & Department of Physics, Fudan University, Shanghai, People's Republic of China⁸ Laboratori Nazionali di Frascati INFN, Frascati (Rome), Italy⁹ INFN sezione Roma 'Tor Vergata', I-00133 Rome, Italy

* Author to whom any correspondence should be addressed.

E-mail: fieldsres@gmail.com**Keywords:** Bayesian active inference, generative model, quantum reference frame, tomographic measurement, topological quantum neural networkSupplementary material for this article is available [online](#)

Abstract

We show how any finite physical system with morphological, i.e. three-dimensional embedding or shape, degrees of freedom and locally limited free energy will, under the constraints of the free energy principle, evolve over time towards a neuromorphic morphology that supports hierarchical computations in which each 'level' of the hierarchy enacts a coarse-graining of its inputs, and dually, a fine-graining of its outputs. Such hierarchies occur throughout biology, from the architectures of intracellular signal transduction pathways to the large-scale organization of perception and action cycles in the mammalian brain. The close formal connections between cone-cocone diagrams (CCCD) as models of quantum reference frames on the one hand, and between CCCDs and topological quantum field theories on the other, allow the representation of such computations in the fully-general quantum-computational framework of topological quantum neural networks.

1. Introduction

The quest to understand how collections of cells form nervous systems that give rise to cognitive capacities has driven research into computational systems using architectures observed in neural tissues. The fundamentals of neuromorphic computing, i.e. computing with systems having functional architectures similar or analogous to those of biological neurons, can be traced back to the work of Mead [1], who pioneered the implementation of very large-scale integration (VLSI) methods. These kinds of functional (neuromimetic) architectures use analog components that approximately mimic neurobiological systems, and were conducive to solving real-world problems with high efficiency and low cost. Hybrid analog-digital systems emulating spiking neurons were also developed as an alternative to purely analog models [2]. Since then, neuromorphic computers have evolved to further emulate the computational architectures of neurons and of functional networks of neurons (for recent reviews, see [3–10]).

As living systems, both neurons and networks of neurons implement computation, in part, using morphology; differential delays between signals, for example, can be implemented by dendritic or axonal processes of different lengths and widths. Changes in morphology also contribute to the implementation of learning; for example, growing or regressing dendritic spines facilitates or inhibits synapse formation and hence location-specific interneural communication [11–14]. Spike-based and structural plasticity together implement memory-write circuits amenable to neuromorphic design [15] (and references therein). At the network scale, activity-dependent pruning during neural development shapes both short- and long-range

cortical connectivity [16–18]. Hence from a biological perspective, a key feature of neuromorphic computing is that it is dynamic: changes in morphology implement changes in computation and vice-versa. This is exemplified in applications of hybrid analog/digital VLSI devices implemented as neuromorphic vision sensors that model concept-learning in relatively simple biological neural networks (BNNs), such as described in [19]¹⁰ Neuromorphic computing foregrounds a separation of temporal scales implicit in natural computation; namely, the distinction between fast inference and slow learning; sometimes considered in the light of ‘dynamics on structure’ [21]. However, on the neuromorphic view, structure itself is dynamic, inheriting from fast inference (e.g. activity-dependent plasticity) and scaffolding inference (e.g. cortical hierarchies and other aspects of functional brain architectures that rest upon synaptic connections) [22]. The use of morphology as a computing resource is not, moreover, unique to neurons. It is an ancient biological strategy, employed by plants, fungi, amoeboid cells of diverse lineages, and even microbial biofilms [23–30]. All of these systems could, therefore, be considered ‘neuromorphic’ computers.

Here, we review and extend previous theoretical work suggesting that any finite physical system capable of employing morphology as a computational resource will, if given a sufficiently informative environment but locally-limited free energy, develop—i.e. evolve in time toward—a ‘neuron-like’ morphology. We show, in particular, that this outcome can be expected on the basis of the free energy principle (FEP, [31–33]) applied to systems with morphological degrees of freedom but locally-limited free energy. Locally-limited free energy restricts local measurements to a few degrees of freedom. Predictive power in this case is maximized if the environment is addressed tomographically. Morphological plasticity is a key enabler of and, when suitably abstracted, a requirement for tomographic measurements. As the FEP is completely scale-free, this result applies at any spatiotemporal scale, and is indeed confirmed by systems from the μm scale of intracellular signaling pathways to the planetary scale of the internet. We suggest on this basis that morphological plasticity, even if merely simulated, is an important resource for neuromorphic computing.

In what follows, we first outline, in section 2, the fundamental role of morphology as a computational resource. We then review, in section 3, the basic ideas underlying the FEP, including the definition of variational free energy (VFE) and its interpretation as Bayesian surprisal, and the key concept of a Markov blanket (MB, [34, 35]) separating a time-persistent system from its environment. Here and in later sections we employ the formalisms of both quantum and classical information theories; however, we make no claims about the role, if any, of long-range coherence in biological computation and neither support nor challenge any particular ‘quantum brain’ models. We show in section 4 how defining the MB of a system defines its environment, and consider the thermodynamics of the MB. We note a critical difference between current artificial computing systems and organisms: the rigid segregation in the former between free-energy exchange with the environment (via a power supply and heat exchangers) and data exchange with the environment (via I/O interfaces or APIs). We then consider MBs as measurement surfaces with locally-limited free energy resources in section 5, and show how morphological degrees of freedom enable varying the correlations between measurement sites in nonuniform environments. This enables us to consider, in section 6, how the FEP drives morphologically-plastic systems—with locally-limited free energy resources—to measure the environment’s state tomographically, using measurements made at different locations to reconstruct a best predictive model of the state. We provide a fully-general physical model of this process in section 7, employing the formalism of topological quantum neural networks (TQNNs, [36]). This shows that TQNNs provide general models of neuromorphic systems. We conclude with implications, predictions, and next steps in section 8. As we employ formalism and concepts from several disciplines, a glossary of terms is provided in supplementary material, appendix A.

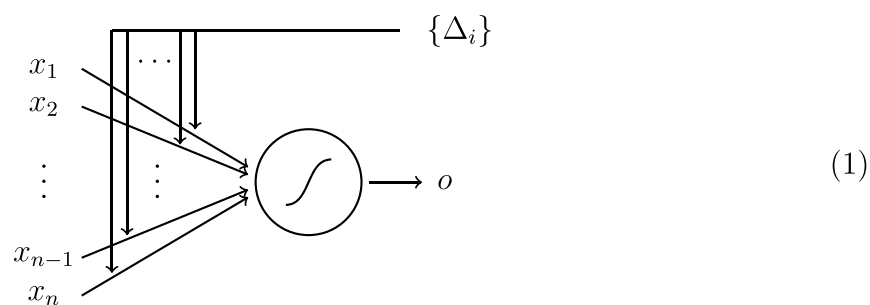
2. Morphology as a computational resource

Abstract models of computation, e.g. the lambda calculus [37] or the Turing machine [38], make no mention of morphology. In conventional computing systems, the three-dimensional (3d) layout or ‘shape’ of the computer is chosen to minimize transmission delays and maximize heat dissipation capacity. Biological systems, in contrast, possess evolved morphologies that determine, in part, how they are able to interact with their worlds. The morphological structures of even single, free-living cells such as amoeba or paramecia encode both intergenerational memory and immediate capabilities for action [39]. The key role of morphology as a computational resource in robotics has been emphasized by Brooks [40] and by the ‘embodied cognition’ movement more generally (see [41] for review).

¹⁰ As reported in [19], the common honeybee stands out as an exemplar having remarkable capabilities for conceptualizing and categorizing (the bees having $\approx 10^6$ neurons compared to $\approx 10^{11}$ neurons in the human brain); in particular, their ability to distinguish between ‘odd’ and ‘even’ numeric quantities [20].

Implementation in 3d space, and hence morphology, imposes two fundamental constraints on computing systems that abstract models ignore. The first is thermodynamic: any irreversible encoding of classical information—writing a bit value to a memory for later retrieval—has a finite free energy cost, at least $\ln 2 k_B T$, where k_B is Boltzmann’s constant and T is ambient temperature [42–44]. This free energy must be obtained from the environment, via a power supply, photosynthesis, or metabolism. The second constraint is on informational coupling: the implementation determines how inputs are obtained from the environment and how outputs are transferred to it. It determines what features or aspects of the environment the computing system can detect—or in psychological language, perceive—and similarly, what features or aspects of the environment it can directly affect by its actions. It is the response to this second constraint that most strongly distinguishes ordinary computers from robots.

The effects of morphology become obvious when comparing conventional artificial neural networks (ANNs) to networks of biological neurons. Setting thermodynamics aside, the ‘environment’ with which an ANN interacts comprises sets of training and test data. It is commonplace to think of an ANN as interacting with, for example, images that have a spatial (here 2d) structure. This, however, is an anthropomorphism; the ANN in fact interacts with sets of finitely-encoded and therefore rational numbers. We can consider a node in a layered, feedforward ANN to have the following structure:



where here $\{x_i\}$ is the set of input values from upstream nodes, $\{\Delta_i\}$ is the set of training (backpropagated error) values, and the rational number o is the output. The ‘sensed environment’ of this node is the ordered pair $(\{x_i\}, \{\Delta_i\})$; the ‘acted-upon environment’ of the node is the rational number o . We will see in section 3 below that these sensed and acted-upon sectors of the environment can be represented formally as sectors of the node’s MB.

Note that *drawing* the node as diagram (1) imposes on it a ‘morphological’ degree of freedom, namely its layout on the 2d Euclidean surface of the page. This, in turn, imposes orders onto the sets $\{x_i\}$ and $\{\Delta_i\}$, making them vectors with the obvious metric. This morphological degree of freedom is not, however, intrinsic to the node; it appears nowhere in a mathematical specification of the function that the node computes, nor does it characterize the (completely abstract) sets $\{x_i\}$ or $\{\Delta_i\}$ or the number o . This *absence of morphology* is, more than absence of hierarchical structure or spiking (which biological neurons can lack), what renders a node in an ANN non-neuromorphic. Nodes in ANNs are non-neuromorphic because they are *amorphic*; they have no morphology.

Implementing an ANN in hardware gives it a morphology: the 3d morphology of the hardware. It also confers a resource requirement for thermodynamic free energy; hence it adds a thermodynamic sector to its environment, which in section 3 below will become a thermodynamic sector of its MB. This exposure to *energetic* exchange with the environment renders the implemented ANN a ‘thing’ in the language of the FEP. How the implemented ANN behaves, i.e. how it regulates its energetic exchange with its environment, determines whether it will persist over time. This regulation of energy exchange in service of persistence, or survival, is the core meaning of embodiment. The ever-present possibility of dysregulation is what renders embodiment ‘precarious’ [45].

Let us now consider a biological neuron, which is by definition embodied and therefore has a morphology. A neuron’s sensed environment is, like the sensed environment of any other system, defined by the sensory structures that it deploys. In the case of a neuron, these are mostly post-synaptic specializations, including clusters of post-synaptic receptors and channels as depicted in [46], figures 4(a) and (b); we will focus on these at the expense of more uniformly distributed biochemical and bioelectric sensors. The neuron’s sensed environment is then the set $\{s_i\}$ of activations detected by these post-synaptic specializations. The neuron’s acted-upon environment is, similarly, the set $\{a_i\}$ of activations generated by its pre-synaptic specializations, again ignoring more uniformly-distributed pumps, secretory systems, etc. These sets $\{s_i\}$ and $\{a_i\}$ comprise the neuron’s MB. Perception and action are linked together by the dynamics on the internal states, which are supported by all internal degrees of freedom of the cell, including

genome, mitochondria and other organelles, cytoskeletal network, etc; these internal dynamics implement the cell's generative model. Hence while it is commonplace to think of a neuron as 'detecting pressure' or 'exciting a muscle' these descriptions are possible only from a larger, tissue-scale perspective. From the neuron's own perspective, it is acting to regulate the bioelectrochemical gradients it detects as state variations of $\{s_i\}$. See [47] for a worked example of dendritic self organization, in terms of structure learning, using the minimization of VFE to implement model selection in terms of dendritic spines. In this example, the morphology of the dendritic tree aligns itself with the temporal sequence of presynaptic inputs that itself depends upon morphology of the neuropil [48]. However, at no point does the (synthetic) neuron 'know' its morphology.

Both inputs to and outputs from a neuron are organized spatially by its morphology. However, the neuron itself cannot detect or represent its morphology, though local changes in morphology are locally detectable, e.g. by differential strain on the cytoskeleton. Hence the neuron's inputs and outputs remain, for the neuron itself, only sets without structure. The overall *function* of the neuron, and hence the values of its outputs, depend however on its computational (i.e. message-passing) architecture and hence on its morphology. It is the role of morphology in determining function—and hence action on the environment—that the FEP explains [46].

The thermodynamic constraints faced by neurons—and indeed on any implemented computing system—force them to trade off the energetic requirements of obtaining data against those of processing data. As discussed in section 6 below, biological systems respond to this tradeoff by coarse-graining, i.e. by compressing their outputs into fewer bits than employed for their inputs. Diagram (1) illustrates this in an extreme case. Coarse-graining leads naturally to a hierarchical computational architecture, and hence to a hierarchical morphology. The combination of high fan-in and hierarchical morphology is spectacularly evident in mammalian cortical neurons, with their elaborate, layer-specific dendritic trees and upwards of 10 000 input synapses [49, 50].

Our goal in what follows is to understand the role of morphology—and in particular, of neuromorphic morphology—as a resource for computation from first principles. To do this, we adopt the formal framework of the FEP, which is applicable, in principle, to any physical system at any scale. After introducing the FEP in the next section, we focus in sections 4–6 on how the structure and functions of the MB surrounding any system define its interaction with its environment and hence shape its morphology. We then show in section 7 how these effects of morphology can be captured in terms of fundamental physical theory.

3. The FEP as a general physical principle

Since its application to brain function [31, 51–53], the variational FEP has been extended into an explanatory framework for living systems at all scales [32, 54–57], along with an extensive scope of related clinical studies (e.g. [58–61]). When formulated as a general principle of classical physics, it characterizes the behavior of all random dynamical systems that remain measurable, and hence identifiable as distinct, persistent entities, over macroscopic times [33]. To summarize, it is shown in [33] that any system that has a non-equilibrium steady state (NESS) solution to its density dynamics (a) possesses an internal dynamics that is conditionally independent of the dynamics of its environment, and (b) will continuously 'self-evidence' by returning its state to (the vicinity of) its NESS. Condition (a) can be thought of as a precondition for any system to have a 'state' that is clearly distinct from the state of its environment; it is effectively the requirement that system and environment are weakly coupled. Given weak coupling and local interactions, the joint system–environment state space can be partitioned into internal (i.e. system), external (i.e. environment) and intermediary MB states. The MB surrounding a set $\{\mu\}$ of internal states comprises all states that are parents of states in $\{\mu\}$, children of states in $\{\mu\}$, or non- $\{\mu\}$ parents of the children of states in $\{\mu\}$, where as usual, the 'parents' and 'children' of a state μ are states with causal arrows into and out of μ , respectively [34, 35]. The MB states can, in turn, be partitioned into sensory states that mediate the influence of external states on internal states and active states that mediate the influence of internal states on external states. In the language of perceptual psychology, the MB functions as an 'interface' [62] that encodes perceptions and actions. It is worth emphasizing that the MB states are elements of the joint system–environment state space; while the MB states are embedded in a physically-continuous spatial boundary in canonical examples such as biological cells, this is not required by the definition of an MB. With this partitioning, Condition (b) then requires that the system behaves so as to preserve the functional integrity of its MB, i.e. that its dynamics does not diverge following a perturbation. The FEP is the statement that any measurable, i.e. bounded and macroscopically persistent, system will behave so as to satisfy these requirements.

More formally, the FEP is a variational or least-action principle stating that a system enclosed by an MB, and therefore having internal states $\mu(t)$ that are conditionally independent of the states $\eta(t)$ of its environment, will evolve in a way that tends to minimize an VFE that is an upper bound on (Bayesian)

surprisal. This free energy is effectively the divergence between the variational density encoded by internal states and the density over external states conditioned on the MB states. If π is a ‘particular’ state $\pi = (b, \mu)$, where $b(t)$ is the state of the MB, the VFE $F(\pi)$ can be written [33, equation (2.3)],

$$\begin{aligned}
 F(\pi) &= \underbrace{\mathbb{E}_{q(\eta)}[\ln q_\mu(\eta) - \ln p(\eta, b)]}_{\text{Variational free energy}} \\
 &= \underbrace{\mathbb{E}_q[-\ln p(b|\eta) - \ln p(\eta)]}_{\text{Energy constraint (likelihood \& prior)}} - \underbrace{\mathbb{E}_q[-\ln q_\mu(\eta)]}_{\text{Entropy}} \\
 &= \underbrace{D_{\text{KL}}[q_\mu(\eta) \| p(\eta)]}_{\text{Complexity}} - \underbrace{\mathbb{E}_q[\ln p(b|\eta)]}_{\text{Accuracy}} \\
 &= \underbrace{D_{\text{KL}}[q_\mu(\eta) \| p(\eta|b)]}_{\text{Divergence}} - \underbrace{\ln p(b)}_{\text{Log evidence}} \geq -\ln p(b).
 \end{aligned} \tag{2}$$

The VFE functional $F(\pi)$ is an upper bound on surprisal (a.k.a. self-information) $\mathfrak{I}(\pi) = -\log P(\pi) = -\ln p(b)$ because the Kullback–Leibler divergence term (D_{KL}) is always non-negative. This KL divergence is between the density over external states η , given the MB state b , and a variational density $Q_\mu(\eta)$ over external states parameterized by the internal state μ . If we view the internal state μ as encoding a posterior over the external state η , minimizing VFE is, effectively, minimizing a prediction error, under a generative model supplied by the NESS density. In this treatment, the NESS density becomes a probabilistic specification of the relationship between external or environmental states and particular (i.e. ‘self’) states. We can interpret the internal and active states in terms of active inference, i.e. a Bayesian mechanics [63], in which their expected flow can be read as perception and action, respectively. In other words, active inference is a process of Bayesian belief updating that incorporates active exploration of the environment. It is one way of interpreting a generalized synchrony between two random dynamical systems that are coupled via an MB.

We have recently reformulated the FEP within a scale-free, spacetime background-free quantum information theory [64]. Quantum information theory provides a particularly simple and convenient representation of physical interaction as information exchange; supplementary material, appendix B provides a brief review. In this representation, the MB is implemented by a decompositional boundary in the joint system-environment Hilbert space that functions as a holographic screen, a topological generalization [65, 66] of the original geometric construction [67–69]. The criterion of conditional independence is implemented by the quantum-theoretic notion of joint-state separability, i.e. absence of entanglement across the holographic screen. The action of the internal system dynamics implements a quantum computation, which can be decomposed as a hierarchy of quantum reference frames (QRFs, [70, 71]). Decomposition into QRFs has the advantage of assigning an explicit semantics, interpretable as a system of units of measurement, to each ‘thread’ of the computation. It is this semantic information that renders measurement outcomes comparable across instances of measurement, and hence renders them ‘differences that make a difference’ [72, 73], that is, differences that are actionable. Each QRF can, in turn, be given a functional specification as a category-theoretic structure, a ‘cone-cocone diagram’ (CCCD) of Barwise–Seligman [74] classifiers. Such CCCDs specify semantically-interpreted information flows, where the semantics are given by the satisfaction conditions of the classifiers, within distributed systems (reviewed in [75, 76]; see also supplementary material, appendix C). In informational/logical terms a CCCD specifies ‘measurement’ and ‘preparation’ as dual memory read/write operations. We have employed this representation to characterize neurons as hierarchical measurement devices [46] as discussed further in section 4 below. Didactically, this allows one to think of perception and action in terms of measurement (read) and preparation (write) operators that stand in for answers (outputs) and questions (inputs), about or from the environment as discussed below.

In either classical or quantum formulations, the FEP provides a generic theory of self-organization for physical systems with sufficient dynamical stability to be identified over time and subjected to multiple measurements, i.e. systems that can be considered ‘things’ that are distinct from their surrounding environments (see especially the discussion of this point in [33]). The MB of any such ‘thing’ underwrites its conditional independence—between its internal states and the external states of its environment—by localizing and thereby restricting information exchange between them. Hence the FEP provides a generic characterization of physical interaction as information exchange, and a generic characterization of internal system dynamics as (Bayesian) inference [33, 63, 64].

This reading of self organization—sometimes referred to as self evidencing [77]—rests, in a foundational way, on the notion of a generative model. Technically, this generative model can be associated with the NESS density over the particular partition of systemic states described above. This density can be factorized into a likelihood (the density over particular states, given their causes; i.e. external states beyond the MB) and a

prior density over particular states that are characteristic of the particle or ‘thing’ in question. The states constitute the attracting set that underwrites the NESS solution to density dynamics. In short, if there exists an MB—defined in terms of conditional dependencies under an NESS density—then there is a lawful description of systemic dynamics that can be cast as gradient flow, asymptotically toward the NESS [33], on a free energy functional of a generative model. Teleologically, the generative model specifies the states to which self-organization (i.e. evidencing) are attracted; namely, the characteristic or preferred states of the ‘thing’ in question. The role of a generative model will be foregrounded in what follows; simply because the structure of a generative model underwrites the dynamics and message-passing we associate with self organization.

4. Defining the MB defines the environment

4.1. Informative versus uninformative sectors

The partitioning of ‘everything’ into ‘system’ and ‘environment’ (where in equation (2) the MB is considered part of the system) built into the FEP formalism has the immediate consequence that every system, by definition, interacts with exactly one other system, its environment. The formalism is, moreover, completely symmetric: the system maintains a well-defined, conditionally-independent state if and only if its environment does as well. We can, indeed, think of system and environment as comprising a generative adversarial network, with each side adapting, as its resources allow, to the other’s actions [78]. This symmetry is particularly manifest in the quantum formalism, which is a completely general representation of two systems (i.e. components of a bipartite Hilbert-space decomposition) open to interaction exclusively with each other.

This exclusive coupling of system to environment has two consequences, both of which have been explored more explicitly within the quantum formalism [64]; see also [65, 79, 80] for further discussion. First, all (thermodynamic) free energy acquired by the system from, and all waste heat dissipated by the system to, its environment must traverse the MB. The MB (or in the quantum formulation, the holographic screen \mathcal{B}), is thus partitioned into ‘informative’ (or ‘observed’) and ‘uninformative’ (or ‘unobserved’) sectors as shown in figure 1. The function of the uninformative sector is purely thermodynamic; formally, it exchanges the free energy required to support irreversible classical computation [42–44].

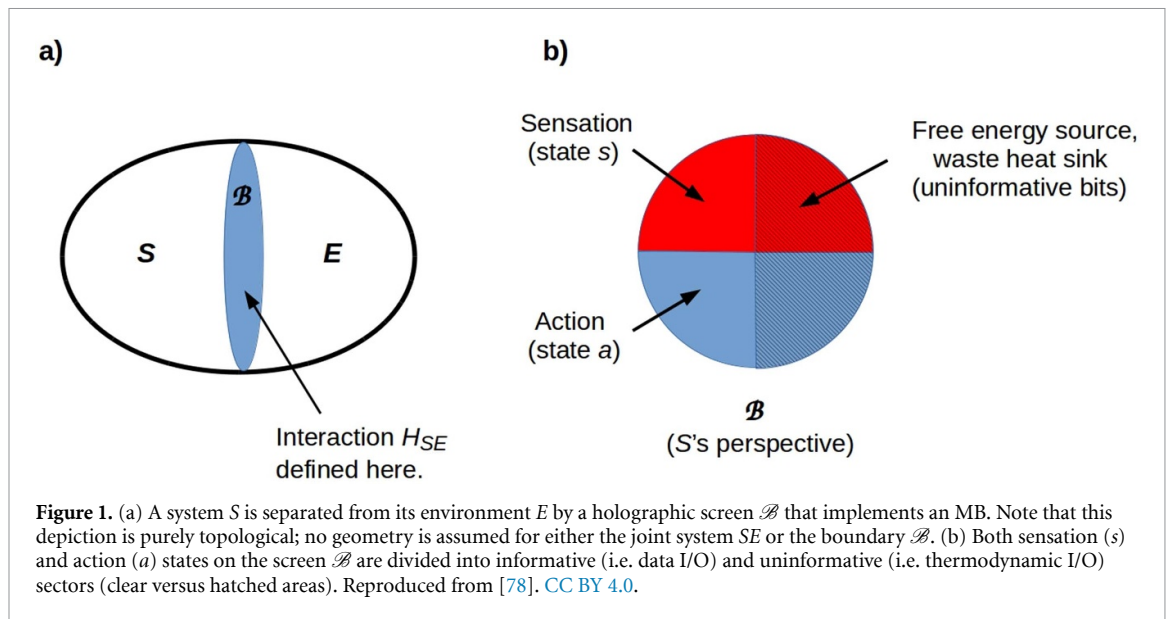
The second consequence of the system–environment decomposition is that the system of interest S has no access to the decompositional, or in quantum terms entanglement, structure of its environment E . Any ‘objects’ detected by S in E are in fact sectors of mutually correlated components of the state of the MB, or in the quantum formulation, sectors of mutually correlated bits encoded on the screen \mathcal{B} [64, 65, 79, 80]. The informative sector of the MB can, therefore, be thought of as implementing an applications programming interface (API) between S and E . Read and write operations to this API are implemented by the internal dynamics of S (respectively, E). In the quantum formulation, these are implemented by QRFs that effectively define the ‘data structures’ encoded on each face of \mathcal{B} .

4.2. Learning is learning a message-passing structure

On the classical view of the Bayesian mechanics entailed by the FEP, the minimization of VFE can be usefully considered at different timescales. For example, optimizing the states or activities of a BNN or ANN is distinct from optimizing the connections or weights; which is distinct from optimising the structure of the neural network per se. These three aspects of VFE minimization map neatly to the distinction between inference, learning and model selection (a.k.a., structure learning), respectively. We start with this observation because, to anticipate the discussion in section 4 below, the structure just is the computational architecture in question and thereby specifies the nature of the message-passing entailed by inference and learning. Morphology in 3d physical space is an implementation resource for computational architecture, as 3d layout in VLSI exemplifies.

On this view, the structure or morphology of any ‘thing’ is subject to the same imperatives as the message-passing; namely to maximize morphological or model evidence (or minimize the associated VFE bound). This can be cast as structure learning in radical constructivism [81–83], or Bayesian model selection in statistics [84]. The implication here is that any morphology must be a ‘good’ model of how its sensory states are caused by external states. This is just an expression of the good regulator theorem from early formulations of self organization in cybernetics [85, 86]. In other words, statistical correlations beyond the MB must be installed in the generative model, in terms of sparse coupling (i.e. message-passing) among internal states (which themselves are ‘things’ equipped with MBs). So what kind of structures, architectures or morphology might one expect to find in things that are good models of their external milieu?

If morphology maximizes model evidence, then any morphology effectively encodes a model specifying expectations about the environment. Cell membranes, for example, can be viewed as encoding expectations about viscosity and ambient chemical potentials, and skeletal systems can be viewed as encoding



expectations about gravity and the buoyancy of media such as air or water. These morphology-encoded models should comply with Occam's principle—or Jaynes maximum entropy principle [87, 88]—in virtue of having minimal complexity. This follows from the fact that log evidence (i.e. negative surprisal) is accuracy minus complexity. Equation (2) shows that complexity is the degree of belief updating incurred by message-passing. Technically, complexity is the KL divergence between posterior and prior, before and after belief updating. In short, a 'good' model is that which provides an accurate account but is as simple as possible. In turn, this requires the right kind of 'coarse graining' or compression [89], to provide an accurate explanation for impressions on the sensory part of the holographic screen implemented by the MB. So, what kind of coarse graining might emerge in a Universe that features probabilistic structure?

This explanation can only be in terms of 'things' and their lawful relationships as described below. At this point, one can conjecture that things—and the (space-time) background that describes their relationships in a parsimonious fashion—would feature in the structure of generative models or morphology. This is evinced in a compelling way by neuroanatomy, which speaks to a distinction between 'what', 'where' and 'when' in carving the sensorium at its joints. For example, one of the most celebrated aspects of brain connectivity is the separation of dorsal and ventral streams that are thought to encode 'where' and 'what' attributes of visual objects, respectively [90]. The argument here is that knowing 'what' something is does not tell you 'where' it is and vice versa. This statistical independence translates into a morphological separation between the dorsal and ventral streams. This separation minimizes complexity and thereby maximizes the efficiency of (variational) measure-message passing and belief updating in terms of statistical, algorithmic and thermodynamic complexity costs [54, 91]. Similar arguments can be made for a separation of 'what' and 'when' [92]; in the sense that knowing 'what' something is, does not tell you 'when' it was 'there'.

This kind of coarse graining (c.f., carving nature at its joints) is ubiquitous in statistics and physics, where it emerges in the guise of mean field approximations; namely, factorizing a probability density into conditionally independent factors [22, 93–96]. Indeed, VFE and message passing are defined under a mean field approximation to a posterior density [91, 97]. Another important structural or morphological feature of 'good' generative models is their deep or hierarchical structure, with an implicit separation of scales in the genesis of—or explanation for—sensory impressions.

The common theme here is a morphology underwritten by the sparsity or absence of message-passing on some factor graph. This foregrounds the imperatives for shielding or sequestering various internal states from other internal states, which brings us back to MBs; however, these are internal MBs that define an internal morphology or message-passing structure. One might conjecture that much of biological self-organization is concerned with isolation and shielding, as a necessary part of internal autopoiesis (e.g. the role of enzymes and catalysts, gap junctions, and many other highly controllable mechanisms for setting up signaling paths and boundaries [98–101]. This occurs at all scales, from subcellular organelles that partition biophysical and chemical reactions to nascent organ compartment boundaries, to the dynamics that guide which members of a swarm pass messages to which others [102–106].

In this respect, morphogenetic self-organization, seen as a pattern formation, requires each individual cell (and/or its progeny) to occupy its own place in the final morphology, and autopoietic self-assembly

results only when each cell successfully detects local patterning signals as predicted by its own generative model [57, 107]. Morphological development thus implies a pre-determined patterning to which a cell ensemble converges—the so-called Target Morphology [108, 109]. Note that this is an essentially classical statement; it assumes the existence of effectively-classical boundaries and hence distinctions between cells. If each cell minimizes VFE then it infers its correct location and its function within the ensemble [57, 107, 110]. Examples in the case of neurons include assortative neuronal migration towards groups with very close or identical node degrees [111, 112] and amalgamation of groups with a common stimulus, following which they ‘cast a vote’ to decide on how to proceed collectively [113]. More generally, the Good Regulator Theorem again applies when each cell, by evidencing its own existence, can vouch inferentially for the same model as the one of the local group in which it is accommodated. In this way, the cell contributes to the eventual release of effective signaling by the ensemble to other formations. This is a basis for a theoretical framework of autopoiesis expressed in terms of VFE minimization, and hence active inference [107].

A final consideration—afforded by the classical FEP—is that the same Bayesian mechanics must apply in a scale-free fashion [33, 110, 114, 115]. In other words, MBs of MBs (i.e. things composed of things) must evince the same kind of message-passing. For example, the intracellular components of a single cell must have the right morphology to maintain the cell’s MB (e.g. a cell surface). Similarly, the ensemble of cells that constitute a multicellular structure must be so structured to maintain the MB of the tissue or organ in question (e.g. a somatic cell on its endothelial surface) [116, 117]. In a similar vein, this implies that the message-passing between MBs (e.g. cells and organs through to conspecifics and cultures) must (look from the outside as if they) comply with the same free energy minimizing imperatives. This translates into efficient communication at the level of intracellular communication, through to languages with minimal algorithmic complexity. In short, message-passing between ‘things’ should incur the minimum amount of belief updating, while communicating as accurately as possible. In what follows, we will see these themes re-emerge, both in terms of biological intelligence and quantum information theory. The semi-classical limit of a TQNN model, in particular, constructs generalizations with the shortest possible trajectories and the maximum topological information as discussed in section 7.

4.3. Object identification by QRFs

From the perspective of an observer S , ‘things’ are located in the environment E . As is obvious from the definition of an MB, however, S cannot ‘see’ E ; S can only detect encodings on its MB \mathcal{B} . A ‘thing’ for S is, therefore, a cluster of bits on \mathcal{B} with high mutual information and hence high joint predictability. Recognizing a ‘thing’—determining that some bits have high mutual information—requires multiple measurements. In particular, any ‘thing’ X can be considered to have two components, a ‘reference’ component R that maintains a constant (up to measurement resolution and relevant coarse-graining) state (or state density or expectation value), and a ‘pointer’ component P with a time-varying state that S considers ‘the state of interest’ of X [64, 65, 79, 80]. Ordinary items of laboratory apparatus provide a canonical example; one can only identify a voltmeter or an oscilloscope if most of their state variables—size, shape, brand name, etc—remain fixed while the ‘pointer’ variables vary to indicate some measured value [118].

Measurements of the states of R and P can, without loss of generality, be regarded as implemented by QRFs [64, 65, 79, 80]. A QRF is simply a physical system with which a measurement is enacted; such a system is a *quantum* reference frame because, being physical, it must at some suitable scale be regarded as a quantum system, and at that scale it encodes unmeasurable, and hence unencodable or ‘nonfungible’ [71] quantum phase information. Such systems are intrinsically semantic: they report not just values, but also units of measurement that render such values mutually comparable. Even a non-standardized QRF such as the length of one’s arm defines a unit of measurement, although an idiosyncratic one. Hence *repeated* observations, which must determine at minimum the state of R and are therefore measurements, are *intrinsically semantic*: they are actions on the world that yield mutually-comparable, and hence actionable observational outcomes.

In a quantum theoretic formulation, measurement and its dual, state preparation, have the same formal representation; a ‘preparation’ process is just a measurement reversed in time. A QRF is, therefore, a preparation device as well as a measurement device: one can prepare a 0.75 m board with a meter stick, just as one can measure a 0.75 m board. Preparation is an *action* on the environment; preparation and measurement together constitute *interaction*. Indeed any physical interaction can be considered a sequence of alternating preparation and measurement steps, as shown in detail in [64]. This duality is preserved in the classical formulation, but remains implicit (i.e. perception as time-reversed action and vice versa). As we have pointed out, the dual character of preparation and measurement as enacted by QRFs allows their representation, in full generality, by category-theoretic structures, namely, the CCCDs Barwise–Seligman classifiers [74], as constructed in [75, 76]; formal definitions and examples are given in supplementary material, appendix C. This representation has been extensively applied in computer science as reviewed in [75]; we prove its generality in the present setting in [76], to which we refer for details. Such structures have the form:

$$\begin{array}{ccccccc}
\mathcal{A}_1 & \xleftarrow{g_{21}} & \mathcal{A}_2 & \xleftarrow{g_{32}} & \dots & \mathcal{A}_k \\
& & \uparrow h_2 & & & \\
& & \mathbf{C}' & & & \\
& & \downarrow h_1 & & & \\
& & \mathcal{A}_1 & \xleftarrow{g_{21}} & \mathcal{A}_2 & \xleftarrow{g_{32}} & \dots & \mathcal{A}_k \\
& & \uparrow f_2 & & & \\
& & \mathbf{C}' & & & \\
& & \downarrow f_1 & & & \\
& & \mathcal{A}_1 & \xleftarrow{g_{21}} & \mathcal{A}_2 & \xleftarrow{g_{32}} & \dots & \mathcal{A}_k
\end{array}
\quad (3)$$

to the macroscopic [123]. Systems can be characterized by ‘cognitive light cones’ that specify their sensory, action, and memory capacities in each such ‘space’ that they inhabit [117, 124, 125].

When the ‘perceived environment’ of a system is understood in terms of QRF-identified sectors of the system’s MB, it becomes obvious that the entire state of the MB, including its thermodynamic sector, provides the ‘context’ of any observation. The ubiquity of context effects has been emphasized by Dzhamarov *et al*, where it forms the basis of the ‘contextuality by default’ approach to statistical analysis of empirical data [126–129]. Conceptually similar, but employing different methods, is the work of Abramsky *et al* [130–133]. Classical and quantum (i.e. entanglement-based) context effects, both of which can be observed even at the macroscopic scale of human behavior [134, 135], have been rigorously distinguished to a degree of generality by commutativity relations between QRFs [76, 119]. Context effects cross the boundaries between ‘spaces’ detected by distinct QRFs; for example, biochemical and bioelectric effects cross-modulate each other not only in neurons, but in all cell types across phylogeny [136–140].

Systems with morphological degrees of freedom, including macromolecules, cells, multicellular organisms, and multi-organism communities, as well as many artifacts, can employ these degrees of freedom to segregate QRFs from each other and hence to regulate context effects. Folded proteins, for example, can separate distinct binding sites in 3d space, and bacteria can segregate chemotactic receptors from flagellar-motor proteins. Such segregation allows interacting with different parts of the environment in different ways. We focus below on characterizing morphological degrees of freedom, and examining the consequences of QRF segregation for computational architecture.

5. MBs with morphological degrees of freedom

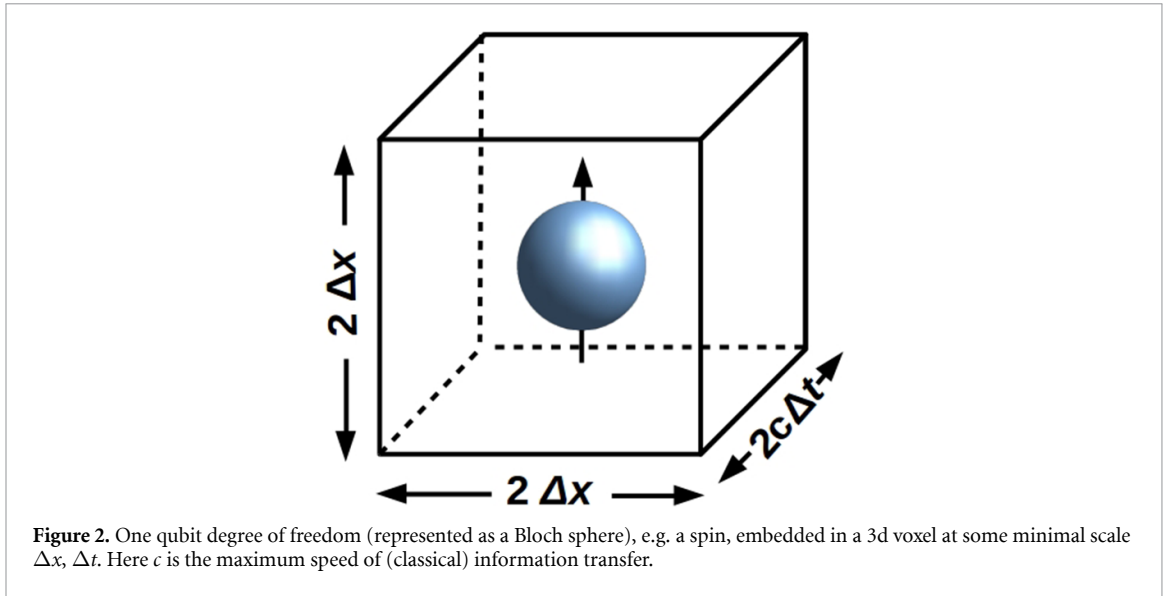
The state b of the MB, or of the screen \mathcal{B} , and in particular the state (s, a) of its informative sector, has thus far been considered a state in some arbitrary (e.g. Hilbert) state space. In particular, no positional (e.g. ordinary Euclidean 3d spatial) degrees of freedom have been assumed. We now add to the MB states, as a parameter, an ancillary ‘morphological’ degree of freedom ξ that, as we will see, is naturally interpretable as a spatial degree of freedom. This ancillary degree of freedom is ancillary in the sense of having no effect on the total system–environment information exchange across the boundary \mathcal{B} ; in the purely-topological notation of figure 1, the interaction H_{SE} does not depend on ξ . As we show below, however, S ’s QRFs partition \mathcal{B} into sectors that are ‘localized’ in the space defined by ξ . Hence ξ usefully parameterizes H_{SE} in a way that a neuron can take advantage of by varying its morphology to selectively deploy its QRFs to specific sectors of \mathcal{B} . We can, therefore, regard H_{SE} as depending locally on ξ ; this will be made explicit in section 7 below when we assign spatial coordinates to the input states of a TQNN.

We further assume that the states $|\xi\rangle$ of ξ (here adopting the Dirac notation for states) are vectors and hence provide a distance measure $\langle\xi|\xi'\rangle$. This effectively ‘geometrizes’ the states b by assigning to each a ‘location’ ξ and allowing ‘distances’ between states to be calculated. In this way ξ plays the role of the 2d geometry of the page in diagram (1); it allows the states b to be placed in an ordered array with the dimension of ξ . From a physical perspective, the simplest geometrization of \mathcal{B} (embedded in a 3d space) represents the space of MB states b as a 2d array of qubits (it hence considers a minimal binary encoding of the states b and implements each bit with a quantum bit, e.g. a spin degree of freedom), and positions each qubit in a voxel of volume $2\Delta x \times 2\Delta x \times 2c\Delta t$ as shown in figure 2, where Δx is the minimal ‘grain size’ of space, Δt is the minimal time to encode one bit, and c is the maximal speed of ‘causal’ classical information transfer. If Δx and Δt are the Planck length and time, respectively, this reproduces the idea of a ‘stretched horizon’ subject to the original, geometric holographic principle, which encodes information at the maximum density given by the Bekenstein bound [67–69]. For biological systems at temperature $T \sim 310$ K, $\Delta t \sim 50$ fs and $\Delta x \sim 1$ Å, and c is the speed of bond-vibration waves in macromolecules [141], a scale roughly 25 orders of magnitude larger than the minimum set by quantum theory.

In order to model neurons, we will assume that ξ has an ‘embedding’ dimension in addition to the ‘2 + 1’ space + time structure shown in figure 2. As we will see in section 6.3 below, the interpretation of this extra dimension depends on the QRFs available to measure it.

Our two principal assumptions can now be stated:

- The state (s, a) of the informative sector of the MB/screen \mathcal{B} is non-uniform in ξ . Parameterizing H_{SE} with ξ therefore reveals local structure in H_{SE} .
- The free energy available via the uninformative sector of the MB/screen \mathcal{B} is sufficiently limited so that only a ‘few’ cycles of classical computation can be performed on each bit in the informative sector.



We will, for simplicity, also assume that the state of the uninformative sector has two components, each with a uniform state. This allows us to treat thermodynamic exchange with the environment as an interaction with ‘hot’ and ‘cold’ heat baths that supply free energy and exhaust waste heat, respectively.

Qualitatively, assumption (a) assures that the informative sector of E is potentially interesting to S , i.e. not perceived merely as noise, while assumption (b) limits S ’s ability to ‘make sense’ of E by performing predictive computations. These assumptions thus keep S in a regime of finite VFE, avoiding the ‘perfect prediction’ limit of $VFE \rightarrow 0$ that, as we show in [64], corresponds to loss of separability (i.e. to an approach to quantum entanglement) between S and E .

On a classical view, these assumptions express the FEP in terms of maximizing accuracy (assumption (a)), while minimizing the complexity cost of belief updating (assumption (b)). In machine learning, a failure to minimize complexity leads to overfitting [142] that can be read as the perfect prediction limit (i.e. quantum entanglement)—topological features of TQNNs, connected to topological quantum field theory (TQFT), have been advocated [143] to explain generalization by DDNs, avoiding both underfitting and overfitting.

We can now develop our main result, showing in section 6 below that given limited free energy, the FEP imposes a hierarchy on the structure of computation over the informative state (s, a) that, when parameterized by the morphological degree of freedom ξ , becomes a tomographical computation defined over effectively ‘spatial’ measurements and producing effectively ‘spatial’ actions. This computation represents the informative sector of E as comprising ‘objects’ that interact ‘causally’ against a spatially-extended background \tilde{E} of non-objects. We then show in section 7 that this action of the FEP can be captured in full generality in the formalism of TQNNs, making explicit that ‘space’ is emergent from connection topology. In this formalism, the role of the spatial embedding (i.e. of ξ) is to enforce a coarse-graining in which the ‘objects’ detected by S are separable and hence statistically conditionally independent. This is exactly the role of ‘space’ in quantum field theories [66]. It allows the mean-field assumption that allows us, as discussed above (section 4.2) in the classical setting, to talk about objects as persistent entities with their own MBs.

6. Tomographic measurements minimize VFE

6.1. ‘Objects’ as sectors in E

We have previously demonstrated the converse of our desired result: that if the bits encoded on a sector $X = RP$ of \mathcal{B} have sufficient mutual information to satisfy the logical criteria (e.g. as encoded by a CCCD) implemented by some QRF $\mathbf{X} = \mathbf{R}\mathbf{P}$ over some macroscopic time interval τ —and in particular, if the measured state density ρ_R of R remains fixed throughout τ —then X will appear to the observer S to be an ‘object’ or ‘persistent thing’ during τ . In particular, the state $|RP\rangle$ (or density ρ_{RP}) will appear to be decoherent from (i.e. not entangled with and hence conditionally independent from) the remainder of \mathcal{B} [64, 65, 79, 80]. Decoherence corresponds, classically, to conditional independence [144, 145]. It is what makes ‘thingness’ and behavioral predictability possible.

Assumption (a) above states that E contains potentially-detectable objects; stated more carefully, assumption (a) states that \mathcal{B} includes sectors that encode bits with significant mutual information. These sectors present S with opportunities for predictability, i.e. for local (on \mathcal{B}) reduction of VFE. Our question is then: what computational (i.e. message-passing) structure can take advantage of ‘islands’ of predictability to minimize VFE while remaining within the free-energy constraint imposed by assumption (b)?

6.2. Hierarchical measurements optimize the accuracy/complexity cost tradeoff

While pure quantum (i.e. unitary) computation costs no free energy, Landauer’s Principle imposes a finite cost on classical bit erasure and hence on classical memory updating [42–44] as noted above. Assumption (b), therefore, effectively limits the writing of classical memories. Recording previous measurement outcomes for comparison with future ones is, therefore, the energetically-limited step in computing predictions. Markov kernels with rational matrix elements provide, for a given (finite) measurement resolution, the most efficient representation of prior measurement outcomes and hence of prior probability distributions [64]. Hence the fundamental energetic tradeoff faced by any VFE-minimizing system—that is, any system compliant with the FEP—is a tradeoff between the resolution with which both prior and posterior probabilities are encoded and the predictive power that they provide¹¹.

The optimal solution to the above tradeoff is, obviously, to *only* encode probabilities that actually contribute to predictive power. Hence we can expect the FEP to drive systems toward identifying and processing input data from only those sectors of their MBs that encode bits with high mutual information, i.e. high redundancy or high error-correction capacity; we will see this also for TQNNs below. Biologically, this corresponds to the (phylogenetic) evolution or (ontogenetic) development of systems with sensory structures and processing pathways specialized to the detectable affordances of their ecological niches [146, 147]. Indeed, many authors have cast natural selection as, implicitly or explicitly, a VFE minimizing process in terms of natural Bayesian model selection or structure learning [148–153].

Sectors of high mutual information induce a connection topology on \mathcal{B} , with these sectors as the open sets. This topological structure breaks the exchange symmetry of bit ‘positions’ (i.e., values of ξ) on \mathcal{B} . This symmetry breaking corresponds to a choice of basis for the Hamiltonian H_{SE} ; again see [64, 65, 79, 80] for details. Under the FEP, the local action of the internal dynamics H_S on each such sector implements a QRF that alternately measures and acts on the bits encoded by that sector. The action of this QRF can, without loss of generality [76], be specified by a CCCD as shown in figure 3.

A limit/colimit and infomorphisms from/to it exist over any mutually-commuting subset of the CCCDs specifying actions of QRFs on \mathcal{B} [141, theorem 7.1]. Hence any mutually-commuting subset of the CCCDs can be hierarchically composed as components of a larger CCCD that processes their combined outputs, as shown in diagram (4). This larger CCCD specifies the action of a QRF that can be thought of as alternately measuring and preparing the states of the component QRFs in the hierarchy. Indeed, this larger QRF induces a single TQFT ([154]) on the collection of sectors measured/prepared by the component QRFs [76]; we will pursue the consequences of this in section 7 below.

Whenever any of the component CCCDs specify (the component QRFs implement) nonlinear processes, e.g., logical AND or XOR, hierarchical decomposition implements coarse-graining. The FEP will drive any system toward such coarse-graining provided the loss of predictive power at the component level is compensated for by a gain of predictive power at the higher level. This will be the case whenever the sectors spanned by the combined CCCD/QRF encode significant mutual information about each other, i.e., in any situation in which there are information *relations* between sectors and hence apparent ‘interactions’ between the ‘objects’ that the sectors represent to S . The FEP, in other words, will drive any system to discover ‘macro variables’ that characterize and ‘emergent causality’ [155, 156] between its identified sectors.

Teleologically speaking, while scientists have only recently developed reliable tools with which to quantitatively track the causal power of different levels of a system [157–164], biological life forms emerging under realistic temporal and energy (metabolic) constraints have always faced selection pressure to estimate causality of meso-scale ‘objects’ in their *Umwelten*. It is essential for survival that an agent spends its precious resources attempting to affect or communicate with (or track) the features of its environment that make a difference, which requires them to coarse-grain their experience into models in which the massive stream of

¹¹ Basically, for any sector X defined by a QRF \mathbf{X} , a generic (k)-time-stamped quantum system A confronts the task of minimizing prediction error $Er_X(k)$ given by

$$Er_X(k) = d(\mathbb{M}_X^A(k), \mathbb{M}_X(k))$$

where $\mathbb{M}_X^A(k)$ and $\mathbb{M}_X(k)$ denote Markov kernels derived from observables, and d is the metric distance between kernels. The FEP in this case asserts that a generic quantum system will act so as to minimize Er_X for each deployable QRF \mathbf{X} (for details, see [64, sections 3 and 4]).

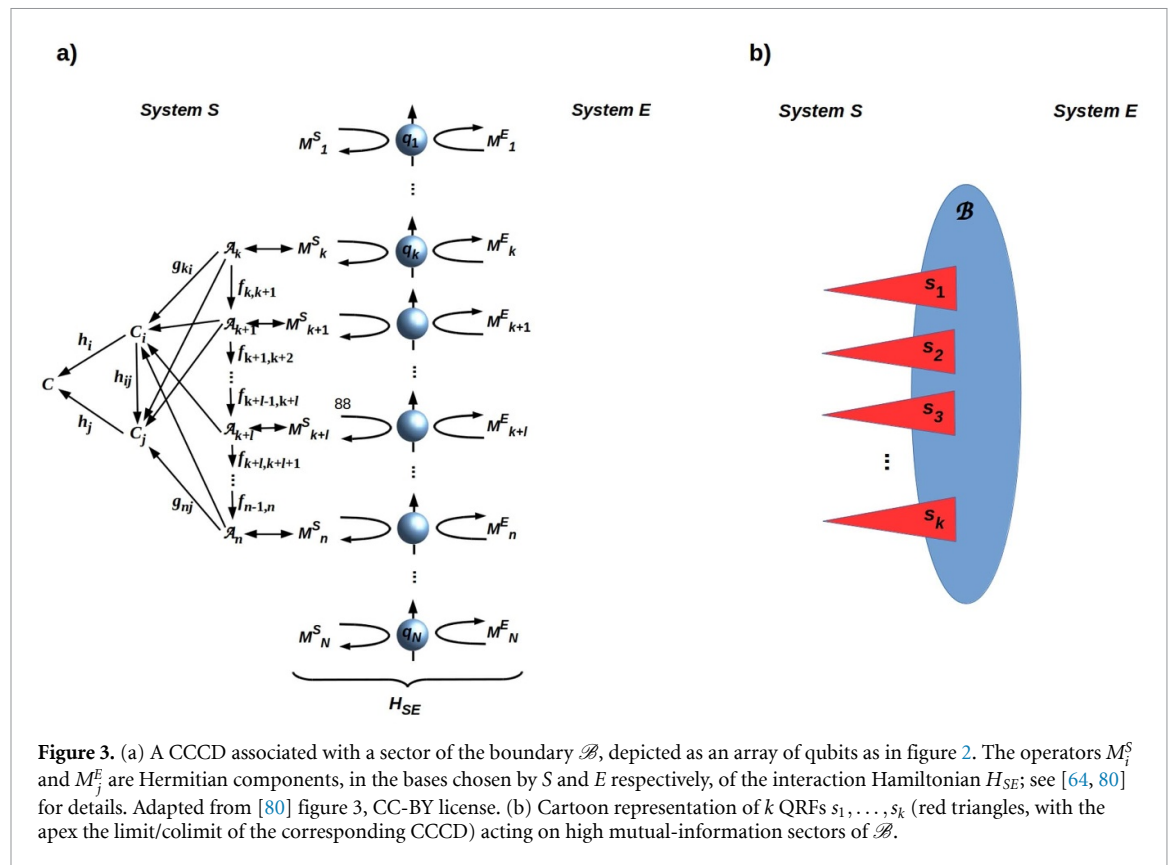


Figure 3. (a) A CCCD associated with a sector of the boundary \mathcal{B} , depicted as an array of qubits as in figure 2. The operators M_i^S and M_j^E are Hermitian components, in the bases chosen by S and E respectively, of the interaction Hamiltonian H_{SE} ; see [64, 80] for details. Adapted from [80] figure 3, CC-BY license. (b) Cartoon representation of k QRFs s_1, \dots, s_k (red triangles, with the apex the limit/colimit of the corresponding CCCD) acting on high mutual-information sectors of \mathcal{B} .

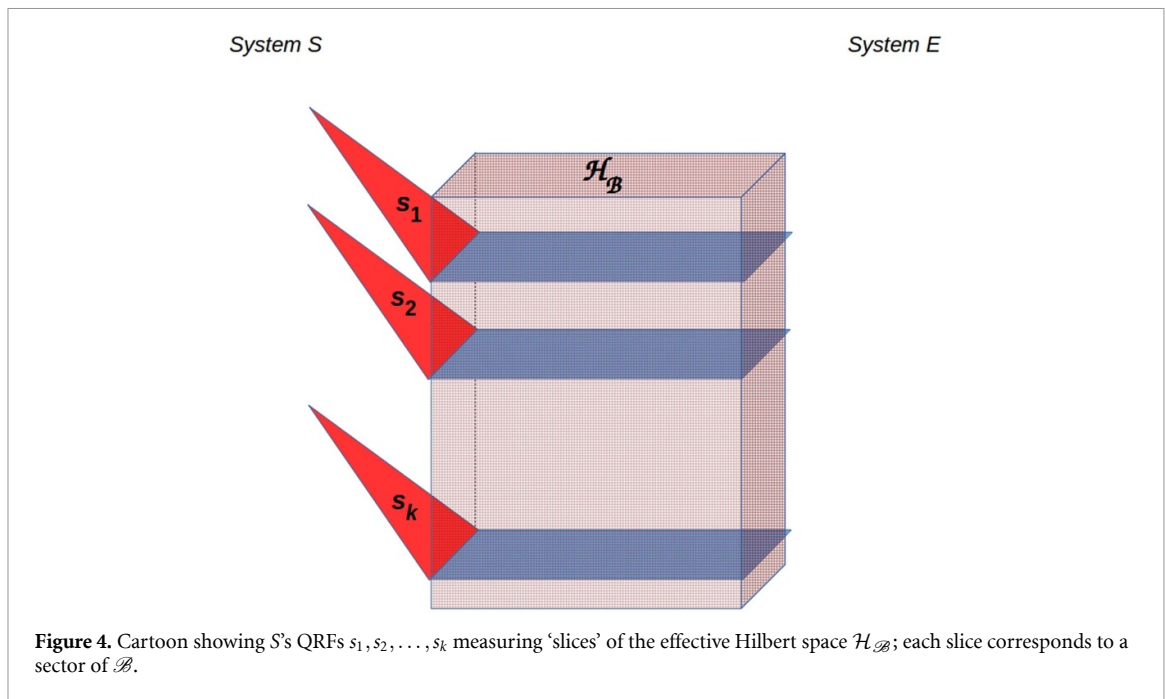
sensory information and potential activities (at all scales) is cut up into convenient ‘objects that do things’ for the purposes of efficient action. Living beings cannot afford a ‘Laplace’s daemon’ (micro-reductionist) view of cause and effect, and the pressure to form models that acknowledge causal potency of higher levels is baked in from the very beginning of the evolution of life.

Biological spiking (e.g., mammalian cortical) neurons are canonical examples of such hierarchical, coarse-graining measurement/preparation systems, with dendritic trees implementing hierarchical measurement and axonal branches implementing hierarchical preparation, namely, action on the surrounding environment [46]. Convolution of post-synaptic potentials at dendritic branch points implement nonlinearities including logical AND and XOR [165]. Gating of action potentials implements similar nonlinearities. These functions are, indeed, precisely the features that most distinguish neurons from simplified models such as diagram (1), and precisely the features that most neuromorphic computing models seek to replicate or at least emulate [3, 4].

6.3. Hierarchical QRFs as tomographic computers

As discussed in section 4.2 above, the relationship between computational structure and morphology exemplified by neurons should characterize any system subject to the FEP. We are now in a position to see why. When the morphological degree of freedom ξ is given the structure of a vector space, it provides a distance measure $\langle \xi | \xi' \rangle$ as discussed in section 4 above; however, we have given no interpretation of this distance. Hierarchical decomposition provides such an interpretation: if sectors of high mutual information are regarded as ‘locations’ on \mathcal{B} —and hence open sets in the connection topology are regarded as simply-connected ‘areas’ in the induced geometry—then mutual information between sectors provides a natural distance measure. A hierarchy of QRFs, in this case, induces a hierarchy of distances on \mathcal{B} that are encoded by the values of the morphological degree of freedom ξ .

Note that the distance measure $\langle \xi | \xi' \rangle$ is a formal description of \mathcal{B} applicable to any interaction H_{SE} in which at least one of the interacting systems, which we by convention label S , has an internal dynamics H_S of sufficient complexity to implement a QRF hierarchy. This does not imply that S itself is capable of measuring the distance $\langle \xi | \xi' \rangle$. Indeed we have explicitly assumed that H_{SE} does not depend on ξ . The system S will be capable of perceiving ‘space’ and thus of measuring distance only if it implements a QRF for spatial measurements. While vertebrates, cephalopods, and some arthropods appear capable of perceiving space, this is not necessary for acting in space (as perceived by us), and the vast majority of organisms, including perhaps all unicells, may lack spatial QRFs altogether [116].



Let us consider what a system S implementing hierarchies of QRFs, but having no *spatial* QRFs perceives. The bits encoded by a high mutual-information sector s_i of the informative sector (s, a) of \mathcal{B} are written by the action of E on that sector. We can think of them, therefore, as outcomes of measurements of basis vectors of the effective Hilbert space $\mathcal{H}_{\mathcal{B}}$ of the MB \mathcal{B} ; this is depicted in figure 3, with the M_i^s as single-qubit (e.g., z -spin s_z) measurement operators. Hence, we can view each of S 's QRFs as measuring a projection or 'slice' of $\mathcal{H}_{\mathcal{B}}$ as shown in figure 4.

Measurements that (partially) reconstruct the state of some system by measuring independent projections of that system are *tomographic* measurements; the (typically hierarchical) process of reconstructing the total state is tomographic reconstruction. Hence QRFs acting on sectors are implementing tomographic measurements of the state of \mathcal{B} , and hierarchies of QRFs are computing a tomographic reconstruction of the state of \mathcal{B} .

By analogy with tomographic reconstruction from image planes as implemented by Positron Emission Tomography in medical imaging, we can think of these slices as having two spatial dimensions as depicted in figure 4. The 'depth' dimension (horizontal in figure 4) is, in this case, notionally perpendicular to the 2d 'surface' of \mathcal{B} . This depth dimension can, therefore, be identified with the embedding dimension of ξ assumed in section 4 above.

The notion of 'depth' and hence the embedding dimension of ξ has, in addition, a second interpretation in terms of time; it can be identified with the total time required to process the inputs from a particular sector through the multiple layers of a hierarchical QRF. This means we can think of the embedded morphology of S in two complementary ways: the morphology both 'extends into' the state space being measured (i.e., into $\mathcal{H}_{\mathcal{B}}$) and 'wraps around' the hierarchical computational structure, conferring a spatial (or space-like) structure on computational (or message-passing) time. This dual aspect of embedding is evident in neurons, which both extend into their (3d) environments and require more time to process distal inputs than proximal ones. In consequence, distal signals are degraded in time resolution and lower in amplitude, rendering time resolution (relative to some proximal standard) and amplitude (relative to some proximal standard) alternative interpretations of the embedding dimension. As discussed in [46], neurons also perform state tomography in time, measuring multiple temporal replicates of input activity patterns to reconstruct the relatively slowly-varying state of the (individual neuron's) environment as a whole.

Perhaps the most celebrated examples of spatiotemporal encoding in the brain are the characteristic responses of the hippocampus; variously read in terms of encoding space—with place and (hierarchically superordinate) grid cells—or, tellingly, as having a key role in memory and the encoding of time [166–182]. From the current perspective, the very existence of place cells—and perhaps receptive fields more generally—speak exactly to the coarse graining of the brain's implicit explanations for its sensorium. And, crucially, how neuronal computations leverage, or are scaffolded by, the morphology of neurons and the connectomes in which they are embedded. This is manifest at many levels empirically; ranging from the emergence of late waveforms in event-related potential research attributed to deep hierarchical processing

[183]; through to temporal gradients as we pass from the back (visual) brain to the front ascending both hierarchal depth and temporal scales [184–187], and indeed the functional anatomy of the hippocampus [188]. The dénouement of this analysis suggests that space and time are perceptual constructs that supervene on the ordinal structure and scheduling of message-passing on neuronal architectures with a certain morphology—a morphology that is apt to explain sensory exchange with the environment, as it is actively queried through (what we perceive as) navigation [61, 189].

What we have said so far here suggests a rather consequential hypothesis: that the translational, rotational, and boost symmetries of spacetime—hence the Poincaré group and Special Relativity—are consequences, for any observer S , of the independence of the Hilbert space of \mathcal{B} from the ancillary coordinate ξ . Observing these symmetries requires identifying, via specific QRFs, objects as sectors on \mathcal{B} . Symmetries of motion are, therefore, consequences of the conservation by S of object identity through time, something that follows immediately from defining, for each object, a reference component R with fixed ρ_R . Hence they are, effectively, consequences of the structure of H_S as a quantum computation. We will pursue this remarkable hypothesis elsewhere; we turn now to the construction of TQNNs as models of neuromorphic computation.

7. TQNNs as general neuromorphic systems

We show in [76] how any sequence of measurements by some fixed QRF on some fixed sector(s) of a boundary \mathcal{B} induces a TQFT on that sector (or those sectors). The proof is completely general and proceeds by two category-theoretic constructions. First, we show that any QRF can be represented by some CCCD, and construct a category **CCCD** of CCCDs that represent QRFs. The morphisms in this category represent transitions between QRFs and thus represent sequential measurements. We then construct a functor from **CCCD** to the category **Cob** of finite cobordisms. A TQFT is a functor from **Cob** to the category **Hilb** of (in this case finite-dimensional) Hilbert spaces, which interprets each cobordism as a manifold of linear maps between Hilbert spaces that serve as boundaries. Transitions between QRFs correspond, therefore, to manifolds of maps between Hilbert spaces, i.e. to TQFTs; see supplementary material, appendix D for a more detailed sketch of this proof.

A TQFT on a boundary \mathcal{B} can be thought of as encoding all possible smooth transformations (e.g. all possible Feynman paths) from some initial configuration \mathcal{B}_{in} of \mathcal{B} to some final configuration \mathcal{B}_{out} . Similarly, a TQFT induced on some sector of \mathcal{B} by sequential actions of one or more QRFs encodes all possible Feynman paths of that sector. One immediate physical consequence is that all effective field theories on observed sectors must be gauge invariant [66, 76, 190]; see [76, 190] for further physical implications of this construction. With a suitable choice of basis, such TQFTs can be implemented by TQNNs [36, 76, 143]. These generalize conventional quantum ANNs [191, 192] by allowing the number and organization of ‘layers’ to be indeterminate. We start by clarifying this peculiar, novel quantum feature of TQNNs, which stands at the forefront of the implementation of TQFTs in machine learning.

TQNNs implement computations on quantum states of the Hilbert space associated to the boundary sectors \mathcal{B} , and can be expanded on the spin-network basis [36, 76, 143]. Spin-networks are in turn supported on 1-complexes (graphs or loops) embedded on the boundary sectors \mathcal{B} , and are colored by certain irreducible representations (irreps) of whatever symmetry describes the system. Note that this embedding requires \mathcal{B} to have, or be extended to have, spatial (i.e., ξ) degrees of freedom. The relevant symmetry is then either individuated by some Lie group or by some quantum group, namely a non-trivial Hopf-algebra [193, 194] (see e.g. [195] for a review of some general extensions of these topics). TQNNs are then represented as superpositions of the basis elements on the boundary sector \mathcal{B} . Spin-networks provide an orthonormal basis, but is worth reminding that loop states as well span the Hilbert space of quantum states over \mathcal{B} , and that a unitary transform exists [196] that connects the two classes of states. The dynamical evolution of the TQNN is then described by TQFT transition amplitudes between \mathcal{B}_{in} and \mathcal{B}_{out} , which are supported on 2-complexes interpolating the TQNN’s 1-complexes [36, 76]. The role of symmetry, as customary in any (effective) quantum field theory, is crucial in recovering the dynamics, as it dictates the interaction among different TQNNs/quantum states embedded on the boundary sectors \mathcal{B}_{in} and \mathcal{B}_{out} . The superposition principle induces a summation over an infinite set of interpolating 2-complexes, supporting virtually fluctuating TQNNs, namely the sum over all the possible evolutions of colored quantum states. The superposition principle also forces to sum over all the possible colours, i.e., all possible assignments of irreps of the symmetry group appropriate to describe a certain physical data set. Topological changes in the graph’s connectivity can be induced, starting from a maximal graph, by solely considering the sum over the irreps. This is due to the fact that assigning a vanishing irrep to a link of a graph corresponds to removing that link from the graph. Then all the interpolating topologies can be obtained summing over all the compatible irreps and intertwiner quantum numbers assigned respectively to the graph nodes and links.

The minimal number of spatial (i.e., ξ) dimensions required to embed 1-complex (graph and/or loop) degrees of freedom in a boundary manifold \mathcal{B} is 2, but a 2d spatial embedding does not allow linking and knotting (i.e., entangling) of these states to take place, which instead requires at least 3 spatial dimensions. A larger number of Hausdorff dimensions is also achievable, and this would induce an encoding of higher-dimensional topological features depending of the number of dimensions of the ambient space.

As noted in [76], the construction of TQNNs presents several formal analogies with quantum gravity (QG) models. Inspecting the dynamics within the framework of the *BF* formalism, just taking into account *BF* Lagrangians, we can easily convince ourselves that topological *BF* theory can be accomplished: in $(1 + 1)$ -dimensions resorting to the Jackiw–Teitelboim gravity [197, 198], either for a $SO(2, 1)$ or for a $SO(3)$ gauge symmetry, in $(2 + 1)$ -dimensions, considering either the topological Einstein–Hilbert action, without accounting for the cosmological constant, or a double Chern–Simons theory, the quantization of which has been proved in [199] to be equivalent to the Turaev–Viro [200, 201] quantization of the Einstein–Hilbert action with cosmological constant, and in $(3 + 1)$ -dimensions, producing the topological Ooguri [202] and Crane–Yetter [203, 204] models. That different realizations that can be recovered shows that dynamics is affected by the number of dimensions, even when considering simplified topological theories described by kinetic *BF* actions.

Having considered a mathematical framework in which functorial evolutions of graphs are supported on 2-complexes, of which graphs are slices of co-dimension 1, and having embedded both the structures respectively in boundary space sectors \mathcal{B} and bulk spaces, of which \mathcal{B} are slices, it is natural to be convinced that ‘spatially’ organized data, localized on the \mathcal{B} sectors, are analyzed hierarchically, in the space or parameter (respectively, in the Euclidean and Lorentzian signature case) flow that induces the slicing of the bulk. This simple consideration has profound consequences since it ensures that TQNNs are ‘neuromorphic’ in a relevant sense. In general, boundary states can be thought as holographic states embedded in lower dimensional projections of the bulk. Boundary states may then encode information about the local curvature when quantum group irreps are considered. For instance, if the irreps that are considered participate in the evaluation of the partition function of the double Chern–Simons theory in 3-dimensions, i.e., the Turaev–Viro model endowed with $SU_q(2)$ irreps, the cosmological constant provides the curvature of the faces that belong to the polyhedra dual to the lattice structure.

We can cast the previous framework in terms of CCCDs, as models of QRFs. Since data are organized spatially, as quantum boundary states/TQNNs are embedded in auxiliary spaces, the boundary sectors \mathcal{B} , CCCDs automatically turn out to be hierarchical in their representations. This implies an orientation for the convolution of CCCD, the inputs of which, thus, are not all processed by only one combinatoric criterion extended to the whole system. The auxiliary spatial degree of freedom participates in the coarse-graining. Because of the proven hierarchical structure, local correlations turn out to be more informative than distant ones, as correlations are suppressed spatially, in inverse powers of the distances involved in the auxiliary spaces. This is a relevant feature for this framework, and involves a confrontation among possible alternative scenarios for coarse-graining: the renormalization group flow à la Wilson versus the Kadanov group—see e.g. [205]. Within the TQNN framework accounted for here, an invariance under coarse-graining of the simplicial tessellation of the space slices of the manifolds emerges, unless one involves a more refined and theoretically challenging extra geometrical renormalization group flow construction—see e.g. the Ricci flow renormalization group approach of [206]. A paradigmatic example is once again provided by the Turaev–Viro model, which is invariant under refinement of the triangulations. This is not true, for instance, for the Ponzano–Regge model [207], which instantiates the spin-foam/*BF* like quantization of the Einstein Hilbert action.

It is crucial to notice, in order to make our considerations in this paper sharper, that the dynamics of TQNNs is instantiated by quantum curvature constraints proper of TQFTs in a way that is equivalent to imposing the FEP on the system. The imposition at the quantum level of the curvature constraints amounts indeed to an extremization of the classical action of either the TQFTs or the effective QFTs describing the specific systems at hand. Indeed, within the semiclassical limit, a constraint that imposes the limitation of the free energy in computational tasks is automatically recovered, imposing tight requirements to the efficiency of TQNN algorithms. Efficiency, which can be modeled as a cost per link in either the TQNN or the CCCD picture, then corresponds to a path minimization for the two-complexes structures intertwining among the boundary states. A detailed analysis of the link between the FEP and the semiclassical limit of TQNNs, which is relevant to unveil generalization in deep-learning systems (DNNs) [143], will be addressed elsewhere.

Finally, in concluding this section we emphasize that we have established TQNNs as a general framework for neuromorphic computation. Notoriously, this is not the case for standard classical DNNs, which rather constitute a less general framework, corresponding to a specific (semi-classical) limit of TQNNs.

Table 1. Correspondence between features of biological or neuromorphic architectures and the formal or physical constructs employed here.

| Bio/neuromorphic architecture | Formal or physical constructs |
|---|---|
| System-environment boundary | MB, holographic screen \mathcal{B} |
| System identity preserved through time | Separability, conditional independence |
| Homeostasis, allostasis | VFE minimization, asymptotic approach to NESS |
| Metabolism | Free energy transfer across \mathcal{B} |
| Information processing | TQNN(s) implemented by internal dynamics |
| Memory for previous observational outcomes | Write/read QRFs acting on \mathcal{B} |
| Object identification, categorization, object-directed behavior | Implementation of object-specific QRFs |

8. Conclusion

The results reviewed here show how any system with morphological degrees of freedom and locally limited free energy will, under the constraints imposed by the FEP, evolve toward a neuromorphic morphology that supports hierarchical computations in which each ‘level’ of the hierarchy enacts a coarse-graining of its inputs, and dually a fine-graining of its outputs. Such hierarchies occur throughout biology, from the architectures of intracellular signal-transduction pathways to the large-scale organization of perception and action processing in the mammalian brain. The close formal connections between CCCDs as models of QRFs on the one hand, and between CCCDs and TQFTs on the other, allow the representation of such computations in the fully-general quantum-computational framework of TQNNs. The mapping between biological or neuromorphic architectures and the formal or physical constructs employed here is summarized in table 1. As noted in the Introduction, we employ these quantum-theoretic constructs because they are scale-free and completely general, but defer discussion of any particular ‘quantum brain’ type models to future work.

One practical implication of the above analysis—that inherits from the distinction between states and parameters of a generative model—is a fundamental distinction between biomimetic computation on Turing machines and neuromorphic computing. From a classical perspective, optimizing the states of a neural network can be read as inference, while optimizing the model parameters (i.e. connection weights in an ANN) corresponds to learning at a slow timescale. In biomimetic schemes, the connection weights or model parameters are generally stored in working memory in the form of tensors to compute the messages that are passed along nodes of a factor graph to instantiate inference at a fast timescale. However, in practice, the vast majority of compute time (and, thermodynamic expenditure) is taken by reading and writing the connectivity tensors from memory. This means that the arguments based upon minimizing the complexity of generative models only provide a lower bound on the thermodynamics of belief updating [42, 208–211]. This lower bound that can only be realized if the connection weights are physically realized as in neuromorphic architectures. This may be an important motivation that goes beyond biomimetic aspirations [142], especially in applications such as edge computing (e.g. surveillance drones).

A further pragmatic perspective on recent trends in machine learning is afforded by the notion of hierarchical computation. In virtue of the fact that these entail a local minimization of VFE (with locally limited thermodynamic free energy), efficient computing on deep networks should conform to these local constraints. Indeed, this is apparent in the move away from backpropagation schemes to local energy-based schemes [212]. This is nicely illustrated by the comparative analyses of backpropagation with predictive coding implementations of deep learning [213–215]. In the current setting, hierarchical predictive coding can be regarded as an implementation of VFE minimization, under hierarchical generative models [216].

More generally, the above results offer some directions for future research. The first is understanding how the pressures that result in neuromorphic architectures impact evolutionary developmental biology, which seeks to determine the origin of specific nervous system patterns [217–219]. More than looking backwards, however, this kind of work can drive advances in both bio-hybrid (biorobotics, chimeric) and software-based AI. A variety of hybrot, organoids, and biobots are being created [220–222] as a way to escape the fact that all of Earth’s biological forms are basically an $N = 1$ example of evolution (barring advances in exobiology). The inclusion of neural (and non-neural bio-electrical) components in these synthetic beings, often made in the absence of any genetic change [124, 125, 223–225], will help test predictions of generic laws driving the structure and function of the body-wide communication system. Similarly, these principles could be of use in designing unconventional and traditional connectionist computational systems, as well as help drive the discovery of interventions guiding cell behavior in regenerative medicine settings [108].

Data availability statement

No new data were created or analysed in this study.

Acknowledgments

K J F is supported by funding for the Wellcome Centre for Human Neuroimaging (Reference: 205103/Z/16/Z), a Canada-UK Artificial Intelligence Initiative (Reference: ES/T01279X/1) and the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 945539 (Human Brain Project SGA3). M L gratefully acknowledges funding from the Guy Foundation and the Finding Genius Foundation. A M wishes to acknowledge support by the Shanghai Municipality, through the Grant No. KBH1512299, by Fudan University, through the Grant No. JJH1512105, the Natural Science Foundation of China, through the Grant No. 11875113, and by the Department of Physics at Fudan University, through the Grant No. IDH1512092/001.

Conflict of interest

The authors declare no competing, financial, or commercial interests in this research.

ORCID iDs

Chris Fields  <https://orcid.org/0000-0002-4812-0744>

Karl Friston  <https://orcid.org/0000-0001-7984-8909>

James F Glazebrook  <https://orcid.org/0000-0001-8335-221X>

Michael Levin  <https://orcid.org/0000-0001-7292-8084>

Antonino Marcianò  <https://orcid.org/0000-0003-4719-110X>

References

- [1] Mead C 1990 Neuromorphic electronic systems *Proc. IEEE* **78** 1629–36
- [2] DeYong M, Findley R and Fields C 1992 The design, fabrication and test of a new VLSI hybrid analog-digital neural processing element *IEEE Trans. Neural Netw.* **3** 363–74
- [3] Schuman C D, Potok T E, Patton R M, Birdwell D, Dean M E, Rose G S and Plank J S 2017 A survey of neuromorphic computing and neural networks in hardware (arXiv:1705.06963v1 [cs.NE])
- [4] Tang J *et al* 2019 Bridging biological and artificial neural networks with emerging neuromorphic devices: fundamentals, progress and challenges *Adv. Mater.* **31** 1902761
- [5] Yang J-Q, Wang R, Ren Y, Mao J-Y, Wang Z-P, Zhou Y and Han S-T 2020 Neuromorphic engineering: from biological to spike-based hardware nervous systems *Adv. Mater.* **32** 2003610
- [6] Zhu Y, Zhu Y, Mao H, He Y, Jiang S, Zhu L, Chen C, Wan C and Wan Q 2021 Recent advances in emerging neuromorphic computing and perception devices *J. Phys. D: Appl. Phys.* **55** 053002
- [7] Rose G S, Shawkat M S A, Foshie A Z, Murray J J V I and Adnan M M 2021 A system design perspective on neuromorphic computer processors *Neuromorph. Comput. Eng.* **1** 022001
- [8] Marković D, Mizrahi A, Querlioz D and Grollier J 2020 Physics for neuromorphic computing *Nat. Rev. Phys.* **2** 499–510
- [9] Primavera B A and Shainline J M 2021 Considerations for neuromorphic supercomputing in semiconducting and superconducting optoelectronic hardware *Front. Neurosci.* **15** 732368
- [10] Vora H, Kathiria P, Agrawal S and Patel U 2022 Neuromorphic computing: review of architecture, issues, applications and research opportunities *Recent Innovations in Computing (Lecture Notes in Electrical Engineering vol 855)* ed P K Singh, Y Singh, J K Chhabra, Z Illés and C Verma (Singapore: Springer)
- [11] Butz M, Wörgötter F and van Ooyen A 2009 Activity-dependent structural plasticity *Brain Res. Rev.* **60** 287–305
- [12] Carulli D, Foscarin S and Rossi F 2011 Activity-dependent plasticity and gene expression modifications in the adult CNS *Front. Mol. Neurosci.* **4** 50
- [13] Hogan M K, Hamilton G F and Horner P J 2020 Neural stimulation and molecular mechanisms of plasticity and regeneration: a review *Front. Cell. Neurosci.* **14** 271
- [14] Runge K, Cardoso C and de Chevigne A 2020 Dendritic spine plasticity: function and mechanisms *Front. Synaptic Neurosci.* **12** 36
- [15] Indiveri G and Liu S-C 2015 Memory and information processing in neuromorphic systems *Proc. IEEE* **10** 1–17
- [16] Shatz C J 1990 Impulse activity and the patterning of connections during CNS development *Neuron* **5** 745–56
- [17] Rakic P, Bourgeois J P and Goldman-Rakic P S 1994 Synaptic development of the cerebral cortex: implications for learning, memory and mental illness *Prog. Brain Res.* **102** 227–43
- [18] Petanjik Z, Judaš M, Šimić G, Rasin M R, Uylings H B M, Rakic P and Kostović I 2011 Extraordinary neonatal synaptic spines in the human prefrontal cortex *Proc. Natl Acad. Sci. USA* **108** 13281–6
- [19] Sandin F, Khan A I, Dyer A G, Amin A H H M, Indiveri G, Chicca E and Osipov E 2014 Concept learning in neuromorphic vision systems: what can we learn from insects? *J. Softw. Eng. Appl.* **7** 387–95
- [20] Howard S R, Greentree J, Avarguès-Weber A, Garcia J E, Greentree A D and Dyer A G 2022 Numerosity categorization by parity in an insect and simple neural network *Front. Ecol. Evol.* **10** 805305
- [21] Aertsen A M, Gerstein G L, Habib M K and Palm G 1989 Dynamics of neuronal firing correlation: modulation of “effective connectivity” *J. Neurophysiol.* **61** 900–17
- [22] Parr T, Sajid N and Friston K J 2020 Modules or mean-fields? *Entropy* **22** 552

- [23] Nakagaki T, Yamada H and Toth A 2000 Maze-solving by an amoeboid organism *Nature* **407** 470
- [24] Blackiston D J, Casey E S and Weiss M R 2008 Retention of memory through metamorphosis: can a moth remember what it learned as a caterpillar? *PLoS One* **3** e1736
- [25] Baluška F, Lev-Yadun S and Mancuso F 2010 Swarm intelligence in plant roots *Trends Ecol. Evol.* **25** 682–3
- [26] Stal L J 2012 Cyanobacterial mats and stromatolites *Ecology of Cyanobacteria II: Their Diversity in Space and Time* ed B A Whitton (Berlin: Springer) pp 65–125
- [27] Vandenberg L N, Adams D S and Levin M 2012 Normalized shape and location of perturbed craniofacial structures in the *Xenopus* tadpole reveal an innate ability to achieve correct morphology *Dev. Dyn.* **241** 863–78
- [28] Müller V C and Hoffmann M 2017 What is morphological computation? On how the body contributes to cognition and control *Artif. Life* **23** 1–24
- [29] Yokawa K and Baluška F 2018 Sense of space: tactile sense for exploratory behavior of roots *Commun. Integr. Biol.* **11** 1–5
- [30] Murugan N J, Kaltman D H, Jin P H, Chien M, Martinez R, Nguyen C Q, Kane A, Novak R, Ingber D E and Levin M 2021 Mechanosensation mediates long-range spatial decision-making in an aneural organism *Adv. Mater.* **2021** 2008161
- [31] Friston K J 2010 The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* **11** 127–38
- [32] Friston K J 2013 Life as we know it *J. R. Soc. Interface* **10** 20130475
- [33] Friston K J 2019 A free energy principle for a particular physics (arXiv:1906.10184 [q-bio.NC])
- [34] Pearl J 1988 *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (San Mateo, CA: Morgan Kaufmann)
- [35] Clark A 2017 How to knit your own Markov blanket: resisting the second law with metamorphic minds *Philosophy and Predictive Processing* vol 3, ed T Wetzinger and W Wiese (Frankfurt am Mainz: Mind Group) p 19
- [36] Marcianò A, Chen D, Fabrocini F, Fields C, Greco E, Gresnigt N, Jinkub K, Lulli M, Terzidis K and Zappala E 2022 Quantum neural networks and topological quantum field theories *Neural Netw.* **153** 164–78
- [37] Church A 1936 A note on the *entscheidungsproblem* *J. Symbol. Logic* **1** 40–41
- [38] Turing A M 1936 On computable numbers, with an application to the *Entscheidungsproblem* *Proc. London Math. Soc.* **s2-42** 230–65
- [39] Fields C and Levin M 2018 Multiscale memory and bioelectric error correction in the cytoplasm-cytoskeleton-membrane system *Wiley Interdiscip. Rev.-Syst. Biol. Med.* **10** e1410
- [40] Brooks R 1986 A robust layered control system for a mobile robot *IEEE J. Robot. Autom.* **2** 14–23
- [41] Anderson M L 2003 Embodied cognition: a field guide *Artif. Intell.* **149** 91–130
- [42] Landauer R 1961 Irreversibility and heat generation in the computing process *IBM J. Res. Dev.* **5** 183–95
- [43] Landauer R 1999 Information is a physical entity *Physica A* **263** 63–67
- [44] Bennett C H 1982 The thermodynamics of computation *Int. J. Theor. Phys.* **21** 905–40
- [45] Weber A and Varela F 2000 Life after kant: natural purposes and the autopoietic foundations of biological individuality *Phenom. Cogn. Sci.* **1** 97–125
- [46] Fields C, Glazebrook J F and Levin M 2022 Neurons as hierarchies of quantum reference frames *BioSystems* **219** 104714
- [47] Kiebel S J and Friston K J 2011 Free energy and dendritic self-organization *Front. Syst. Neurosci.* **5** 80
- [48] Branco T, Clark B A and Hausser M 2010 Dendritic discrimination of temporal input sequences in cortical neurons *Science* **329** 1671–5
- [49] Spruston N 2008 Pyramidal neurons: Dendritic structure and synaptic integration *Nat. Rev. Neurosci.* **9** 206–21
- [50] Rasia-Filho A A, Guerra K T K, Vásquez C E, Dall'Oglio A, Reberger R, Jung C R and Calcagnotto M E 2021 The subcortical-allocortical-neocortical continuum for the emergence and morphological heterogeneity of pyramidal neurons in the human brain *Front. Synapt. Neurosci.* **13** 616607
- [51] Friston K J 2005 A theory of cortical responses *Phil. Trans. R. Soc. B* **360** 815–36
- [52] Friston K J, Kilner J and Harrison L 2006 A free energy principle for the brain *J. Physiol. Paris* **100** 70–87
- [53] Friston K J and Stephan K E 2007 Free-energy and the brain *Synthese* **159** 417–58
- [54] Friston K J, FitzGerald T, Rigoli F, Schwartenbeck P and Pezzulo G 2017 Active inference: a process theory *Neural Comput.* **29** 1–49
- [55] Ramstead M J, Badcock P B and Friston K J 2018 Answering Schrödinger's question: a free-energy formulation *Phys. Life Rev.* **24** 1–16
- [56] Ramstead M J, Constant A, Badcock P B and Friston K J 2019 Variational ecology and the physics of sentient systems *Phys. Life Rev.* **31** 188–205
- [57] Kuchling F, Friston K, Georgiev G and Levin M 2020 Morphogenesis as Bayesian inference: a variational approach to pattern formation and control in complex biological systems *Phys. Life Rev.* **33** 88–108
- [58] Peters A, McEwen B S and Friston K 2017 Uncertainty and stress: why it causes diseases and how it can be mastered by the brain *Prog. Neurobiol.* **156** 164–88
- [59] Friston K, Brown H R, Siemerikus J and Stephan K E 2016 The dysconnection hypothesis (2016) *Schizophr. Res.* **176** 83–94
- [60] McGovern H T, De Foe A, Biddell H, Leptourgos P, Corlett P, Bandara K and Hutchinson B T 2022 Learned uncertainty: the free energy principle in anxiety *Front. Psychol.* **13** 943785
- [61] Friston K and Buzsáki G 2016 The functional anatomy of time: what and when in the brain *Trends Cogn. Sci.* **20** 500–11
- [62] Hoffman D D, Singh M and Prakash C 2015 The interface theory of perception *Psychon. Bull. Rev.* **22** 1480–506
- [63] Ramstead M J, Sakthivadivel D A R, Heins C, Koudahl M, Millidge B, Da Costa L, Klein B and Friston K J 2022 On Bayesian mechanics: a physics of and by beliefs (arXiv:2205.11543 [cond-mat.stat-mech])
- [64] Fields C, Friston K, Glazebrook J F and Levin M 2021 A free energy principle for generic quantum systems *Prog. Biophys. Mol. Biol.* **173** 36–59
- [65] Fields C and Marcianò A 2019 Holographic screens are classical information channels *Quant. Rep.* **2** 326–36
- [66] Addazi A, Chen P, Fabrocini F, Fields C, Greco E, Lulli M, Marcianò A and Pasechnik R 2021 Generalized holographic principle, gauge invariance and the emergence of gravity à la Wilczek *Front. Astron. Space Sci.* **8** 563450
- [67] 't Hooft G 1993 Dimensional reduction in quantum gravity *Salamfestschrift* ed A Ali, J Ellis and S Randjbar-Daemi (Singapore: World Scientific) pp 284–96
- [68] Susskind L 1995 The world as a hologram *J. Math. Phys.* **36** 6377–96
- [69] Bousso R 2002 The holographic principle *Rev. Mod. Phys.* **74** 825–74
- [70] Aharonov Y and Kaufherr T 1984 Quantum frames of reference *Phys. Rev. D* **30** 368–85
- [71] Bartlett S D, Rudolph T and Spekkens R W 2007 Reference frames, super-selection rules and quantum information *Rev. Mod. Phys.* **79** 555–609

- [72] Bateson G 1972 *Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution and Epistemology* (Northvale, NJ: Jason Aronson)
- [73] Roederer J 2005 *Information and Its Role in Nature* (Berlin: Springer)
- [74] Barwise J and Seligman J 1997 *Information Flow: The Logic of Distributed Systems* (Cambridge Tracts in Theoretical Computer Science vol 44) (Cambridge: Cambridge University Press)
- [75] Fields C and Glazebrook J F 2019 A mosaic of Chu spaces and channel theory I: category-theoretic concepts and tools *J. Exp. Theor. Artif. Intell.* **31** 177–213
- [76] Fields C, Glazebrook J F and Marcianò A 2022 Sequential measurements, topological quantum field theories and topological quantum neural networks *Fortschr. Phys.* **70** 2200104
- [77] Hohwy J 2016 The self-evidencing brain *Noûs* **50** 259–85
- [78] Kuchling F, Fields C and Levin M 2022 Metacognition as a consequence of competing evolutionary time scales *Entropy* **24** 601
- [79] Fields C and Glazebrook J F 2020 Representing measurement as a thermodynamic symmetry breaking *Symmetry* **12** 810
- [80] Fields C, Glazebrook J F and Marcianò A 2021 Reference frame induced symmetry breaking on holographic screens *Symmetry* **13** 408
- [81] Tenenbaum J B, Kemp C, Griffiths T L and Goodman N D 2011 How to grow a mind: statistics, structure and abstraction *Science* **331** 1279–85
- [82] Salakhutdinov R, Tenenbaum J B and Torralba A 2013 Learning with hierarchical-deep models *IEEE Trans. Pattern Anal. Mach. Intell.* **35** 1958–71
- [83] Tervo D G, Tenenbaum J B and Gershman S J 2016 Toward the neural implementation of structure learning *Curr. Opin. Neurobiol.* **37** 99–105
- [84] Friston K and Penny W 2011 Post hoc Bayesian model selection *NeuroImage* **56** 2089–99
- [85] Conant R C and Ashby W R 1970 Every Good Regulator of a system must be a model of that system *Int. J. Syst. Sci.* **1** 89–97
- [86] Seth A K 2014 The cybernetic Bayesian brain: from interoceptive inference to sensorimotor contingencies *Open Mind* ed T Metzinger and J M Windt (Frankfurt am Main: MIND Group) ch 35
- [87] Jaynes E T 1957 Information theory and statistical mechanics *Phys. Rev. II* **106** 620–30
- [88] Sakthivadivel D A R 2022 A constraint geometry for inference and integration (arXiv:2203.08119)
- [89] Schmidhuber J 2010 Formal theory of creativity, fun and intrinsic motivation (1990–2010) *IEEE Trans. Auton. Mental Dev.* **2** 230–47
- [90] Ungerleider L G and Haxby J V 1994 “What” and “where” in the human brain *Curr. Opin. Neurobiol.* **4** 157–65
- [91] Winn J Bishop C M 2005 Variational message passing *J. Mach. Learn. Res.* **6** 661–94
- [92] Friston K and Buzsáki G 2016 The functional anatomy of time: what and when in the brain *Trends Cogn. Sci.* **20** 500–11
- [93] Yedidia J S, Freeman W T and Weiss Y 2005 Constructing free-energy approximations and generalized belief propagation algorithms *IEEE Trans. Inform. Theory* **51** 2282–2312
- [94] Dauwels J 2007 On variational message passing on factor graphs 2007 *IEEE Int. Symp. on Information Theory*
- [95] Zhang C, Butepage J, Kjellström H and Mandt S 2018 Advances in variational inference *IEEE Trans. Pattern Anal. Mach. Intell.* **41** 2008–26
- [96] Parr T, Markovic D, Kiebel S J and Friston K J 2019 Neuronal message passing using mean-field, Bethe and marginal approximations *Sci. Rep.* **9** 1889
- [97] Beal M J 2003 Variational algorithms for approximate Bayesian inference *PhD Thesis* University College London
- [98] Mathews J and Levin M 2017 Gap junctional signaling in pattern regulation: physiological network connectivity instructs growth and form *Dev. Neurobiol.* **77** 643–73
- [99] Yamashita Y M, Inaba M and Buszczak M 2018 Specialized intercellular communications via cytonemes and nanotubes *Annu. Rev. Cell Dev. Biol.* **34** 59–84
- [100] Naphade S, Sharma J, Gaide Chevonnay H P, Shook M A, Yeagy B A, Rocca C J, Ur S N, Lau A J, Courtoy P J and Cherqui S 2015 Brief reports: lysosomal cross-correction by hematopoietic stem cell-derived macrophages via tunneling nanotubes *Stem Cells* **33** 301–9
- [101] Wang X, Veruki M L, Bukoreshtliev N V, Hartveit E and Gerdes H-H 2010 Animal cells connected by nanotubes can be electrically coupled through interposed gap-junction channels *Proc. Natl Acad. Sci. USA* **107** 17194–9
- [102] Turner J S 2011 Termites as models of swarm cognition *Swarm Intell.* **5** 19–43
- [103] Deisboeck T S and Couzin I D 2009 Collective behavior in cancer cell populations *BioEssays* **31** 190–7
- [104] Couzin I D 2009 Collective cognition in animal groups *Trends Cogn. Sci.* **13** 36–43
- [105] Shapiro J A 1995 The significances of bacterial colony patterns *BioEssays* **17** 597–607
- [106] Shapiro J A 1998 Thinking about bacterial populations as multicellular organisms *Annu. Rev. Microbiol.* **52** 81–104
- [107] Friston K, Levin M, Sengupta B and Pezzulo G 2015 Knowing one’s place: a free-energy approach to pattern regulation *J. R. Soc. Interface* **12** 20141383
- [108] Pezzulo G and Levin M 2015 Re-membering the body: applications of computational neuroscience to the top-down control of regeneration of limbs and other complex organs *Integr. Biol.* **7** 1487–517
- [109] Pezzulo G and Levin M 2016 Top-down models in biology: explanation and control of complex living systems above the molecular level *J. R. Soc. Interface* **13** 20160555
- [110] Palacios E R, Razi A, Parr T, Kirchhoff M and Friston K 2020 On Markov blankets and hierarchical self-organization *J. Theor. Biol.* **486** 110089
- [111] Barrat A, Barthélemy M and Vespignani A 2008 *Dynamical Processes on Complex Networks* (Cambridge: Cambridge University Press)
- [112] Rubinov M and Sporns O 2010 Complex network measures of brain connectivity: uses and interpretations *NeuroImage* **52** 1059–69
- [113] Latham P and Dayan P 2005 Touché: the feeling of choice *Nat. Neurosci.* **8** 408–9
- [114] Parr T, Da Costa L and Friston K 2020 Markov blankets, information geometry and stochastic thermodynamics *Phil. Trans. R. Soc. A* **378** 20190159
- [115] Da Costa L et al 2021 Bayesian mechanics for stationary processes *Phil. Trans. R. Soc. A* **477** 20210518
- [116] Fields C, Glazebrook J F and Levin M 2021 Minimal physicalism as a scale-free substrate for cognition and consciousness *Neurosci. Conscious.* **7** niab013
- [117] Levin M 2019 The computational boundary of a “self”: developmental bioelectricity drives multicellularity and scale-free cognition *Front. Psychol.* **10** 1688

- [118] Fields C 2018 Some consequences of the thermodynamic cost of system identification *Entropy* **20** 797
- [119] Fields C and Glazebrook J F 2022 Information flow in context-dependent hierarchical Bayesian inference *J. Exp. Theor. Artif. Intell.* **34** 111–42
- [120] Ashby W R 1956 *Introduction to Cybernetics* (London: Chapman and Hall)
- [121] Moore E F 1956 Gedanken experiments on sequential machines *Autonoma Studies* ed C W Shannon and J McCarthy (Princeton, NJ: Princeton University Press) pp 129–55
- [122] Robbins R J, Krishtalka L and Wooley J C 2016 Advances in biodiversity: metagenomics and the unveiling of biological dark matter *Stand. Genomic Sci.* **11** 69
- [123] Fields C and Levin M 2022 Competency in navigating arbitrary spaces as an invariant for analyzing cognition in diverse embodiments *Entropy* **24** 819
- [124] Levin M 2021 Life, death and self: Fundamental questions of primitive cognition viewed through the lens of body plasticity and synthetic organisms *Biochem. Biophys. Res. Commun.* **564** 114–33
- [125] Levin M 2022 Technological approach to mind everywhere: an experimentally-grounded framework for understanding diverse bodies and minds *Front. Syst. Neurosci.* **16** 768201
- [126] Dzharfarov E N, Kujala J V and Cervantes V H 2016 Contextuality-by-default: a brief overview of concepts and terminology *Quantum Interaction Lecture Notes in Computer Science* vol 9525 ed H Atmanspacher, T Filik and E Pothos (Berlin: Springer) pp 12–23
- [127] Dzharfarov E N and Kujala J V 2017 Contextuality-by-default 2.0: systems with binary random variables *Quantum Interaction (Lecture Notes in Computer Science* vol 10106) ed J A Barros, B Coecke and E Pothos (Berlin: Springer) pp 16–32
- [128] Dzharfarov E N, Cervantes V H and Kujala J V 2017 Contextuality in canonical systems of random variables *Phil. Trans. R. Soc. A* **375** 20160389
- [129] Dzharfarov E N and Kon M 2018 On universality of classical probability with contextually labeled random variables *J. Math. Psychol.* **85** 17–24
- [130] Abramsky S and Brandenburger A 2011 The sheaf-theoretic structure of non-locality and contextuality *New J. Phys.* **13** 113036
- [131] Abramsky S and Brandenburger A 2014 An operational interpretation of negative probabilities and no-signaling models *Panangaden Festschrift (Lecture Notes in Computer Science* vol 8464) ed F van Bruegel et al (Cham: Springer International Publishing) pp 59–74
- [132] Abramsky S, Barbosa R S and Mansfield S 2017 Contextual fraction as a measure of contextuality *Phys. Rev. Lett.* **119** 050504
- [133] Abramsky S and Hardy L 2012 Logical Bell inequalities *Phys. Rev. A* **85** 062114
- [134] Cervantes V H and Dzharfarov E N 2018 Snow Queen is evil and beautiful: experimental evidence for probabilistic contextuality in human choices *Decision* **5** 193–204
- [135] Basieva I, Cervantes V H, Dzharfarov E N and Khrennikov A 2019 True contextuality beats direct influences in human decision making *J. Exp. Psychol. Gen.* **148** 1925–37
- [136] Levin M 2012 Molecular bioelectricity in developmental biology: new tools and recent discoveries *BioEssays* **34** 205–17
- [137] Tseng A and Levin M 2013 Cracking the bioelectric code: probing endogenous ionic controls of pattern formation *Commun. Integr. Biol.* **6** e22595
- [138] Levin M, Pezzulo G and Finkelstein J M 2017 Endogenous bioelectric signaling networks: exploiting voltage gradients for control of growth and form *Annu. Rev. Biomed. Eng.* **19** 353–87
- [139] Levin M and Martyniuk C J 2018 The bioelectric code: an ancient computational medium for dynamic control of growth and form *BioSystems* **164** 76–93
- [140] Fields C, Bischof J and Levin M 2020 Morphological coordination: a common ancestral function unifying neural and non-neural signaling *Physiology* **35** 16–30
- [141] Fields C and Levin M 2021 Metabolic limits on classical information processing by biological cells *BioSystems* **209** 104513
- [142] Sengupta B, Friston K 2018 How robust are deep neural networks? (arXiv:1804.11313)
- [143] Marciano A, Chen D, Fabrocini F, Fields C, Lulli M and Zappala E 2022 Deep neural networks as the semi-classical limit of topological quantum neural networks (arXiv:2210.13741)
- [144] Zurek W H 2003 Decoherence, einselection and the quantum origins of the classical *Rev. Mod. Phys.* **75** 715–75
- [145] Schlosshauer M 2007 *Decoherence and the Quantum to Classical Transition* (Berlin: Springer)
- [146] Fields C and Levin M 2020 Integrating evolutionary and developmental thinking into a scale-free biology *BioEssays* **42** 1900228
- [147] Fields C and Levin M 2020 Does evolution have a target morphology? *Organisms* **4** 57–76
- [148] Campbell J O 2016 Universal Darwinism as a process of Bayesian inference *Front. Syst. Neurosci.* **10** 49
- [149] Ramirez J C and Marshall J A R 2017 Can natural selection encode Bayesian priors? *J. Theor. Biol.* **426** 57–66
- [150] Da Costa L et al 2020 Natural selection finds natural gradient (arXiv:2001.08028)
- [151] Vanchurin V, Wolf Y I, Katsnelson M I and Koonin E V 2022 Toward a theory of evolution as multilevel learning *Proc. Natl Acad. Sci. USA* **119** e2120037119
- [152] Frank S A 2012 Natural selection V. How to read the fundamental equations of evolutionary change in terms of information theory *J. Evol. Biol.* **25** 2377–96
- [153] Sella G and Hirsh A E 2005 The application of statistical physics to evolutionary biology *Proc. Natl Acad. Sci. USA* **102** 9541–6
- [154] Atiyah M 1988 Topological quantum field theory *Pub. Math. IHES* **68** 175–86
- [155] Hoel E P, Albantakis L and Tononi G 2013 Quantifying causal emergence shows that macro can beat micro *Proc. Natl Acad. Sci. USA* **110** 19790–5
- [156] Hoel E P 2017 When the map is better than the territory *Entropy* **19** 188
- [157] Hoel E and Levin M 2020 Emergence of informative higher scales in biological systems: A computational toolkit for optimal prediction and control *Commun. Integr. Biol.* **13** 108–18
- [158] Hoel E P 2018 Agent above, atom below: how agents causally emerge from their underlying microphysics *Wandering Towards a Goal: How Can Mindless Mathematical Laws Give Rise to Aims and Intention?* ed A Aguirre, B Foster and Z Merali (Cham: Springer International Publishing) pp 63–76
- [159] Albantakis L et al 2017 What caused what? An irreducible account of actual causation (arXiv:1708.06716)
- [160] Hoel E P, Albantakis L, Marshall W and Tononi G 2016 Can the macro beat the micro? Integrated information across spatiotemporal scales *Neurosci. Conscious.* **2016** niw012
- [161] Clif O M, Lizier J T, Wang X R, Wang P, Obst O and Prokopenko M 2017 Quantifying long-range interactions and coherent structure in multi-agent dynamics *Artif. Life* **23** 34–57

- [162] Wibral M, Lizier J T, Vögler S, Priesemann V and Galuske R 2014 Local active information storage as a tool to understand distributed neural information processing *Front. Neuroinform.* **8** 1
- [163] Wang X R, Miller J M, Lizier J T, Prokopenko M and Rossi L F 2012 Quantifying and tracing information cascades in swarms *PLoS One* **7** e40084
- [164] Lizier J T, Heinze J, Horstmann A, Haynes J-D and Prokopenko M 2011 Multivariate information-theoretic measures reveal directed information structure and task relevant changes in fMRI connectivity *J. Comput. Neurosci.* **30** 85–107
- [165] Gidon A, Zolnik T A, Fidzinski P, Bolduan F, Papoutsis A, Poirazi P, Holtkamp M, Vida I and Larkum M E 2020 Dendritic action potentials and computation in human layer 2/3 cortical neurons *Science* **367** 83–87
- [166] Milner B, Corkin S and Teuber H L 1968 Further analysis of the hippocampal amnesic syndrome: fourteen year follow-up study of H.M. *Neuropsychologia* **6** 215–34
- [167] O'Keefe J and Dostrovsky J 1971 The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat *Brain Res.* **34** 171–5
- [168] Buzsáki G, Chen L S and Gage F 1990 Spatial organization of physiological activity in the hippocampal region: Relevance to memory formation *Prog. Brain Res.* **83** 257–68
- [169] O'Keefe J and Recce M L 1993 Phase relationship between hippocampal place units and the EEG theta rhythm *Hippocampus* **3** 317–30
- [170] Buzsáki G 1998 Memory consolidation during sleep: a neurophysiological perspective *J. Sleep Res.* **7** 17–23
- [171] Burgess N, Jackson A, Hartley T and O'Keefe J 2000 Predictions derived from modelling the hippocampal role in navigation *Biol. Cybern.* **83** 301–12
- [172] Burgess N, Maguire E A and O'Keefe J 2002 The human hippocampus and spatial and episodic memory *Neuron* **35** 625–41
- [173] Davis M H and Johnsrude I S 2003 Hierarchical processing in spoken language comprehension *J. Neurosci.* **23** 3423–31
- [174] Dragoi G and Buzsáki G 2006 Temporal encoding of place sequences by hippocampal cell assemblies *Neuron* **50** 145–57
- [175] Sejnowski T J and Paulsen O 2006 Network oscillations: emerging computational principles *J. Neurosci.* **26** 1673–6
- [176] Burgess N, Barry C and O'Keefe J 2007 An oscillatory interference model of grid cell firing *Hippocampus* **17** 801–12
- [177] Wittner L, Henze D A, Záborszky L, Buzsáki G, 2007 Three-dimensional reconstruction of the axon arbor of a CA3 pyramidal cell recorded and filled *in vivo* *Brain Struct. Funct.* **212** 75–83
- [178] Moser E I, Kropf E and Moser M B 2008 Place cells, grid cells and the brain's spatial representation system *Annu. Rev. Neurosci.* **31** 69–89
- [179] Buzsáki G and Moser E T 2013 Memory, navigation and theta rhythm in the hippocampal-entorhinal system *Nat. Neurosci.* **16** 130–8
- [180] Bush D, Barry C, Manson D and Burgess N 2015 Using grid cells for navigation *Neuron* **87** 507–20
- [181] Stachenfeld K L, Botvinick M M and Gershman S J 2017 The hippocampus as a predictive map *Nat. Neurosci.* **20** 1643–53
- [182] Barron H C, Auksztulewicz R and Friston K 2020 Prediction and memory: a predictive coding account *Prog. Neurobiol.* **192** 101821
- [183] Garrido M I, Kilner J M, Kiebel S J and Friston K J 2007 Evoked brain responses are generated by feedback loops *Proc. Natl Acad. Sci. USA* **104** 20961–6
- [184] Hasson U, Yang E, Vallines I, Heeger D J and Rubin N 2008 A hierarchy of temporal receptive windows in human cortex *J. Neurosci.* **28** 2539–50
- [185] Kiebel S J, Daunizeau J and Friston K 2008 A hierarchy of time-scales and the brain *PLoS Comput. Biol.* **4** e1000209
- [186] Cocchi L, Sale M V, Gollo L L, Bell P T, Nguyen V T, Zalesky A, Breakspear M and Mattingley J B 2016 A hierarchy of timescales explains distinct effects of local inhibition of primary visual cortex and frontal eye fields *eLife* **5** e15252
- [187] Wang X-J and Kennedy H 2016 Brain structure and dynamics across scales: In search of rules *Curr. Opin. Neurobiol.* **37** 92–98
- [188] Pezzulo G, van der Meer M A A, Lansink C S and Pennartz C M A 2014 Internally generated sequences in learning and executing goal-directed behavior *Trends Cogn. Sci.* **18** 647–57
- [189] Kaplan R and Friston K J 2018 Planning and navigation as active inference *Biol Cybern.* **112** 323–43
- [190] Fields C, Glazebrook J F and Marciano A 2022 The physical meaning of the holographic principle (arXiv:2210.16021 [quant-ph])
- [191] Farhi E and Neven H 2018 Classification with quantum neural networks on near-term processors (arXiv:1802.06002)
- [192] Beer K, Bondarenko D, Farrelly T, Osborne T J, Salzmann R, Scheiermann D and Wolf R 2020 Training deep quantum neural networks *Nat. Commun.* **11** 1–6
- [193] Chari V and Pressley A 1994 *A Guide to Quantum Groups* (Cambridge: Cambridge University Press)
- [194] Majid S 1995 *Foundations of Quantum Group Theory* (Cambridge: Cambridge University Press)
- [195] Baianu I, Glazebrook J F and Brown R 2009 Algebraic topology foundations of supersymmetry and symmetry breaking in quantum field theory and quantum gravity: a review *Symmet. Integra. Geom. Meth. Applic. (SIGMA)* **5** 051
- [196] Rovelli C and Smolin L 1995 Spin networks and quantum gravity *Phys. Rev. D* **52** 5743–59
- [197] Jackiw R 1985 Lower dimensional gravity *Nucl. Phys. B* **252** 343–56
- [198] Teitelboim C 1983 Gravitation and hamiltonian structure in two spacetime dimensions *Phys. Lett. B* **126** 41–45
- [199] Gresnigt N, Marciano A and Zappala A 2022 On the dynamical emergence of $SU_q(2)$ from the regularization of $2 + 1D$ gravity with cosmological constant (arXiv:2201.01726 [gr-qc])
- [200] Turaev V G and Viro O Y 1992 State sum invariants of 3-manifolds and quantum 6j-symbols *Topology* **31** 865–902
- [201] Turaev V G 1994 *Quantum Invariants of Knots and 3-Manifolds* (New York: de Gruyter)
- [202] Ooguri H 1992 Topological lattice models in four dimensions *Mod. Phys. Lett. A* **7** 2799–2810
- [203] Crane L and Yetter D 1993 A categorical construction of 4-D topological quantum field theories *Quantum Topology* ed L Kauffman and R Baadhio (Singapore: World Scientific) pp 120–3
- [204] Crane L, Kauffman L and Yetter D 1997 State-sum invariants of 4-manifolds I *J. Knot Theor. Ramifications* **6** 177–234
- [205] Lulli M, Marciano A and Zappala E 2022 in preparation
- [206] Lulli M, Marciano A and Shan X 2021 Stochastic quantization of general relativity à la Ricci-flow (arXiv:2112.01490 [gr-qc])
- [207] Ponzano G and Regge T 1968 Semiclassical limit of Racah coefficients *Spectroscopy and Group Theoretical Methods in Physics* ed F Bloch (Amsterdam: North-Holland)
- [208] Bennett C H 2003 Notes on Landauer's principle, reversible computation and Maxwell's demon *Stud. Hist. Phil. Sci. B* **34** 501–10
- [209] Jarzynski C 1997 Nonequilibrium equality for free energy differences *Phys. Rev. Lett.* **78** 2690–3
- [210] Crooks G E 2007 Measuring thermodynamic length *Phys. Rev. Lett.* **99** 100602
- [211] Still S, Sivak D A, Bell A J and Crooks G E 2012 Thermodynamics of prediction *Phys. Rev. Lett.* **109** 120604

- [212] Scellier B and Bengio Y 2017 Equilibrium propagation: Bridging the gap between energy-based models and backpropagation *Front. Comput. Neurosci.* **11** 24
- [213] Millidge B, Tschantz A and Buckley C L 2020 Predictive coding approximates backprop along arbitrary computation graphs (arXiv:2006.04182)
- [214] Marino J 2021 Predictive coding, variational autoencoders and biological connections *Neural Comput.* **34** 1–44
- [215] Salvatori T *et al* 2021 Reverse differentiation via predictive coding (arXiv:2103.04689)
- [216] Friston K 2008 Hierarchical models in the brain *PLoS Comput. Biol.* **4** e1000211
- [217] Randel N, Shahidi R, Verasztó C, Bezares-Calderón L A, Schmidt S and Jékely G 2015 Inter-individual stereotypy of the *Platynereis* larval visual connectome *eLife* **4** e08069
- [218] Jékely G, Keijzer F and Godfrey-Smith P 2015 An option space for early neural evolution *Phil. Trans. R. Soc. B* **370** 20150181
- [219] Keijzer F, van Duijn M and Lyon P 2013 What nervous systems do: early evolution, input-output and the skin brain thesis *Adapt. Behav.* **21** 67–85
- [220] Clawson W P and Levin M 2022 Endless forms most beautiful: teleonomy and the bioengineering of chimeric and synthetic organisms *Biol. J. Linnean Soc.* **blac073** accepted
- [221] Sole R, Moses M and Forrest S 2019 Liquid brains, solid brains *Phil. Trans. R. Soc. B* **374** 20190040
- [222] Macia J, Vidiella B and Sole R V 2017 Synthetic associative learning in engineered multicellular consortia *J. R. Soc. Interface* **14** 20170158
- [223] Kriegman S, Blackiston D, Levin M and Bongard J 2021 Kinematic self-replication in reconfigurable organisms *Proc. Natl Acad. Sci. USA* **118** e2112672118
- [224] Blackiston D, Lederer E, Kriegman S, Garnier S, Bongard J and Levin M 2021 A cellular platform for the development of synthetic living machines *Sci. Robot.* **6** eabf1571
- [225] Kriegman S, Blackiston D, Levin M and Bongard J 2020 A scalable pipeline for designing reconfigurable organisms *Proc. Natl Acad. Sci. USA* **117** 1853–9
- [226] Friston K, Parr T and de Vries B 2017 The graphical brain: belief propagation and active inference *Netw. Neurosci.* **1** 381–414
- [227] Harris A K 2018 The need for a concept of shape homeostasis *BioSystems* **173** 65–72