Control flow in active inference systems

Part I:

Classical and quantum formulations of active inference

Chris Fields[a*], Filippo Fabrocini[b,c], Karl Friston[d,e], James F. Glazebrook[f,g], Hananel Hazan[a], Michael Levin[a,h] and Antonino Marcianò[i,j,k]

[a] *Allen Discovery Center at Tufts University, Medford, MA 02155 USA*

[b] *College of Design and Innovation, Tongji University, 281 Fuxin Rd, 200092 Shanghai, CHINA*

[c] *Institute for Computing Applications "Mario Picone", Italy National Research Council, Via dei Taurini, 19, 00185 Rome, ITALY*

[d] *Wellcome Centre for Human Neuroimaging, University College London, London, WC1N 3AR, UK*

[e] *VERSES Research Lab, Los Angeles, CA, 90016 USA*

[f] *Department of Mathematics and Computer Science, Eastern Illinois University, Charleston, IL 61920 USA*

[g] *Adjunct Faculty, Department of Mathematics, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA*

[h] *Wyss Institute for Biologically Inspired Engineering at Harvard University, Boston, MA 02115, USA*

[i] *Center for Field Theory and Particle Physics & Department of Physics Fudan University, Shanghai, CHINA*

[j] *Laboratori Nazionali di Frascati INFN, Frascati (Rome), ITALY*

[k] *INFN sezione Roma "Tor Vergata", I-00133 Rome, ITALY*

April 22, 2023

Abstract

Living systems face both environmental complexity and limited access to free-energy resources. Survival under these conditions requires a control system that can activate, or deploy, available perception and action resources in a context specific way. In this Part I, we introduce the free-energy principle (FEP) and the idea of active inference as Bayesian prediction-error minimization, and show how the control problem arises in active inference systems. We then review classical and quantum formulations of the FEP, with the former being the classical limit of the latter. In the accompanying Part II, we show that when systems are described as executing active inference driven by the FEP, their control flow systems can always be represented as tensor networks (TNs). We show how TNs as control systems can be implemented within the general framework of quantum topological neural networks, and discuss the implications of these results for modeling biological systems at multiple scales.

Keywords

Bayesian mechanics; Dynamic attractor; Free-energy principle; Quantum reference frame; Scale-free model; Topological quantum field theory

1 Introduction

Living things offer remarkable examples of complex, multi-level control policies that guide adaptive function at several scales. At the same time, they are made of components which

*Corresponding author at: Allen Discovery Center at Tufts University, Medford, MA 02155 USA; *E-mail address*: fieldsres@gmail.com

are usually thought of as physical objects obeying simple rules; how can these two perspectives be unified in a rigorous manner? The framework of *active inference* answers this question, by providing a completely general, scale-free formal framework for describing interactions between physical systems in cognitive terms. It is based on the Free Energy Principle (FEP), first introduced in neuroscience [1, 2, 3, 4, 5] before being extended to living systems in general [6, 7, 8, 9] and then to all self-organizing systems [10, 11, 12, 13]. The FEP states that any system that interacts with its environment weakly enough to maintain its identifiability over time 1) has a Markov blanket (MB) that separates its internal states from the states of its environment [14, 15, 16, 17, 18] and 2) behaves over time in a way that asymptotically minimizes a variational free energy (VFE) measured at its MB. Equivalently, the FEP states that any system with a non-equilibrium steady-state (NESS) solution to its density dynamics (and hence an MB) will act so as to maintain its state in the vicinity of its NESS. Any system compliant with the FEP can be described as engaging, at all times, in active inference: a cyclic process in which the system observes its environment, updates its probabilistic "Bayesian beliefs" (i.e., posterior or conditional probability densities) over future behaviors, and acts on its environment so as to test its predictions and gain additional information. The internal dynamics of such a system can be described as inverting a generative model (GM) of its environment that furnishes predictions of the consequences of its actions on its MB.

As a fully-general principle, the FEP applies to all physical systems, not just to behaviorally interesting, plausibly cognitive systems, such as organisms or autonomous robots [10]. Intuitively, behavior is interesting – to external observers and, we can assume, to the behaving system itself – when it is complex, situation-appropriate, and robust in the face of changing environmental conditions. Friston et al. [13] characterize interesting systems as "strange particles", whose internal (i.e., cognitive) states are influenced by their actions only via perceived environmental responses; such systems have to "ask questions" of their environments in order to get answers [19]. Such systems, even bacteria and other basal organisms [20, 21, 22, 23], have multiple ways of observing and acting upon their environments and deploy these resources in context-sensitive ways. In operations-research language, they exhibit situational awareness, i.e., awareness of the context of actions [24], and deploy attention systems to manage the informational, thermodynamic, and metabolic costs of maintaining such awareness [12, 22]. Situational awareness is dependent on both short- and long-term memory, or more technically, on the period of time over which precise [Bayesian] beliefs exist, sometimes referred to as the temporal depth or horizon of the GM [20, 21]. Upper limits can, therefore, be placed on behavioral complexity by examining the capacity and control of memory systems from the cellular scale [25] upwards. Living systems from microbial mats to human societies employ stigmergic memories [22] and hence have "extended minds" [26] in the sense of the literature on embodied, embedded, enactive, extended, and affective (4EA) cognition [27, 28]. Such memories must be both readable and writable; hence any system using them must have dedicated, memory-specific perception–action capabilities.

Any system with multiple perception–action (or stimulus–response) capabilities requires a control system that enables context-guided perception and action and precluding the continuous, simultaneous deployment of all available perception–action capabilities. Such self organization entails the selection of a particular course of action – i.e., policy – from all plausible policies entertained by the system's GM. In the active inference framework, the system's internal states – hence its GM – can be read as encoding posterior probability densities (i.e., Bayesian beliefs) over the causes of its sensory states, including, crucially, its own actions. This leads to the notion of planning and control as inference [29, 30, 31], with the ensuing selection of an action given by the most likely policy. In bacteria such as *E. coli*, for example, mutual inhibition between gene regulatory networks (GRNs) for different metabolic operons permit the expression of specific carbon-source (e.g., sugar) metabolism pathways only when the target carbon source is detected in the environment [32]. The control of foraging behavior via chemotaxis employs a similar, in this case bistable, mechanism [33]. Such mechanisms are active in multicellular morphogenesis, for example, in the head-versus-tail morphology decision in planaria [34]. In the human brain, mutual inhibition between competing visual processing streams is evident in binocular rivalry (switching between distinct scenes presented to left and right eyes) or in the changing interpretations of ambiguous figures such as the Necker cube [35, 36]; similar competitive effects are observed in other sensory pathways [37]. It also characterizes the competitive interaction between the dorsal and ventral attention systems, which implement top-down and bottom-up targeting of sensory resources, respectively [38]. It is invoked at a still larger scale in global workspace models of conscious processing, in which incoming information streams must compete, with each inhibiting the others, for "access to consciousness" [39, 40]. Mutual inhibition creates an energetic barrier that the control system that implements switching must expend free-energy resources to overcome; the controller must not only turn "on" the preferred system, but also turn "off" the inhibition. The required free energy expenditure in turn induces hysteresis and hence the non-linear, winner-takes-all "switch" behavior in the time regime. Such barriers and their temporal consequences persist in more complex control systems whenever two perception–action capabilities are either functionally incompatible or too expensive to deploy simultaneously.

Switching between perception–action capabilities can be regarded, from a theoretical, FEP perspective, as selecting a plausible policy, or plan, supported by the GM. Technically, the probability distribution over policies or plans can be computed from a free energy functional expected under the posterior predictive density over possible outcomes, as described in §2.1 below. The control system that implements the switching process can be considered to employ the GM to predict, or assign a probability distribution to, each perception-action capability (i.e., policy) as a function of context [41, 42]. We can consider the GM to generate probabilistic "beliefs" about the consequences of actions, where here a "belief" is just a mathematically-described structure, e.g., a classical conditional probability density or a quantum state with an assigned amplitude. "Planning" or "control" can, therefore, always be cast as inference – again in the basal sense of computation – implemented by variational message passing or "belief propagation" on a (normal style) factor graph: a graph with nodes corresponding to the factors of a probability distribution and undirected edges corresponding to message-passing channels. Factor graphs can be combined with message passing schemes, with the messages generally corresponding to sufficient statistics of the factors in question, to provide an efficient computation of functions such as marginal densities [43, 44]. Hence one can formalize control – under the FEP – in terms of control as inference, which implies that there is a description of control in terms of message passing on a factor graph. When the GM is over discrete states, this implies a description of control in terms of tensor operators.

Nearly all simulations of planning – under discrete state space GMs – use the factor-graph formalism. Crucially, the structure of the factor graph embodies the structure of the GM and, effectively, the way that any system represents the (apparent causes of) data on its MB; i.e., the way it "carves nature at its joints," into states, objects and categorical features. Under the (classical) FEP, the factors that constitute the nodes of the factor graph correspond to the state-space factorization in a mean field approximation, as used by physicists, or by statisticians to implement variational Bayesian (a.k.a., approximate Bayesian) inference [45]. See [46] for technical details, [47] for an application to the brain, and Supplementary Information, Table 1 for a list of selected applications.

We show in Parts I and II of this paper that control flow in such systems can always be formally described as a tensor network, a factorization of some overall tensor (i.e., high-dimensional matrix) operator into multiple component tensor operators that are pairwise contracted on shared degrees of freedom [48]. In particular, we show that the factorization conditions that allow the construction of a TN are exactly the same as those that allow the identification of distinct, mutually conditionally independent (in quantum terms, decoherent), sets of data on the MB, and hence allow the identification of distinct "objects" or "features" in the environment. This equivalence allows the topological structures of TNs – many of which have been well-characterized in applications of the TN formalism to other domains [48] – to be employed as a classification of control structures in active inference systems; including cells, organisms, and multi-organism communities. It allows, in particular, a principled approach to the question of whether, and to what extent, a cognitive system can impose a decompositional or mereological (i.e., part-whole) structure on its environment. Such structures naturally invoke a notion of locality, and hence of geometry. The geometry of spacetime itself has been described as a particular TN – a multiscale entanglement renormalization ansatz (MERA) [49, 50, 51] – suggesting a deep link between control flow in systems capable of observing spacetime (i.e., capable of implementing internal representations of spacetime) and the deep structure of spacetime as a physical construct.

We begin in this Part I, §2 by analyzing the control-flow problem in three different representations of active inference. First, we employ the classical, statistical formulation of the FEP [10, 11] in §2.1 to describe control flow as implementing discrete, probabilistic transitions between dynamical attractors on a manifold of computational states. We then reformulate the physical interaction in quantum information-theoretic terms in §2.2; in this formulation [12], components of the GM can be considered to be distinct quantum reference frames (QRFs) [52, 53] and represented by hierarchical networks of Barwise-Seligman classifiers [54] as developed in [55, 56, 57, 58]. Control flow then implements discrete transitions between QRFs. The third step, in §2.3, employs the mapping between hierarchies of classifiers and topological quantum field theories (TQFTs) developed in [59]. Here, control flow is implemented by a TQFT, with transition amplitudes given by a path integral.

The second and third of these representations provide formal characterizations of intrinsic (or "quantum") context effects that are consistent with both the sheaf-theoretic treatment of contextuality in [60, 61] and the Contextuality by Default (CbD) approach of [62, 63]; see also the discussion in [57] and [59, §7.2]. The underlying theme is that contextuality arises due to the non-existence of any globally definable (maximally connected) conditional probability distribution across all possible observations (see e.g., [64] for a review from a more general physics perspective). Extending our earlier analysis [57], we discuss reasons to expect that active inference systems will generically exhibit such context effects.

In Part II, we develop a fully-general tensor representation of control flow, and prove that this tensor can be factored into a TN if, and only if, the separability (or conditional statistical independence) conditions needed to identify distinct features of, or objects in, the environment are met. We show how TN architecture allows classification of control flows, and give two illustrative examples. We then discuss several established relationships between TNs and artificial neural network (ANN) architectures, and show how these generalize to topological quantum neural networks [59, 65], of which standard deep-learning (DL) architectures are a classical limit [66]. Having developed these formal results, we turn to implications of these results for biology, and discuss how TN architectures correlate with the observational capabilities of the system being modeled, particularly as regards abilities to detect spatial locality and mereology. We consider how to classify known control pathways in terms of TN architecture and how to employ the TN representation of control flow in experimental design. We conclude by looking forward to how these FEP-based tools can further integrate the physical and life sciences.

2 Formal description of the control problem

2.1 The attractor picture

Let U be a random dynamical system that can be decomposed into subsystems with states $\mu(t)$, $b(t)$, and $\eta(t)$ such that the dependence of the $\mu(t)$ on the $\eta(t)$, and vice-versa, is only via the $b(t)$. In this case, the $b(t)$ form an MB separating the $\mu(t)$ from the $\eta(t)$. We will refer to the $\mu(t)$ as "internal" states, to the $\eta(t)$ as "environment" states, and to the combined $\pi(t) = (b(t), \mu(t))$ as "particular" (or "particle") states [10]. The FEP is a variational or least-action principle stating that any system – that interacts sufficiently weakly with its environment – can be considered to be enclosed by an MB, i.e. any "particle" with states $\pi(t) = (b(t), \mu(t))$, will evolve in a way that tends to minimize a variational free energy (VFE) $F(\pi)$ 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. It can be written [10, Eq. 2.3],

$$
\begin{aligned}
F(\pi) &= \underbrace{\mathbb{E}_{q(\eta)}[\ln q_\mu(\eta) - \ln p(\eta, b)]}_{\text{Variational free energy}} \\
&= \underbrace{\mathbb{E}_q[-\ln p(b|\eta) - \ln p(\eta)]}_{\text{Energy constraint (likelihood \& prior)}} - \underbrace{\mathbb{E}_q[-\ln q_\mu(\eta)]}_{\text{Entropy}} \\
&= \underbrace{D_{KL}[q_\mu(\eta)|p(\eta)]}_{\text{Complexity}} - \underbrace{\mathbb{E}_q[\ln p(b|\eta)]}_{\text{Accuracy}} \\
&= \underbrace{D_{KL}[q_\mu(\eta)||p(\eta|b)]}_{\text{Divergence}} \underbrace{-\ln p(b)}_{\text{Log evidence}} \geq -\ln p(b)
\end{aligned}
\tag{1}
$$

The VFE functional $F(\pi)$ is an upper bound on surprisal (a.k.a. self-information) $\mathfrak{I}(\pi) = -\ln p(\pi) > -\ln p(b)$ 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 GM encoded 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 MB states in terms of active inference, i.e., a Bayesian mechanics [11], in which their expected flow can be read as perception and action, respectively. Here "active" states are a subset of the MB states that are not influenced by environmental states and – for the kinds of particles considered here – do not influence internal states. 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.

If the "particle" π is a biological cell, it is natural to consider the MB b to be implemented by the cell membrane and the "internal" states μ to be the internal macromolecular or biochemical states of the cell; indeed, it is this association that motivated the application of the FEP to cellular life [5]. In this case, the NESS corresponds to the state, or neighborhood of states, that maintain homeostasis (or more broadly, allostasis [67, 68, 69]) and hence maintain the structural and functional integrity of π as a living cell. This activity of self-maintenance has been termed "self-evidencing" [70]; systems compliant with the FEP can be considered to be continually generating evidence of – or for – their continued existence [10].

In the terminology of [13] cells are "strange particles" – their signal transduction pathways monitor (components of) the states of their environments, but do not directly monitor their actions on their environments (i.e., their own active states). The consequences of any action can only, therefore, be deduced from the response of the environment. In this situation, causation is always uncertain: whether an action by the environment on the cell – what the cell detects as an environmental state change – is a causal consequence of an action the cell has taken in the past cannot be determined by the data available to the cell. Every action, therefore, increases VFE, while every observation (potentially) decreases it. The (apparent) task of the cell's GM is to minimize the increases, on average, while maximizing the decreases.

The Bayesian mechanics afforded by the FEP implies a (classical) thermodynamics; indeed, the FEP can be read as a constrained maximum entropy or caliber principle [71, 72]. This follows from the fact that inference, i.e., self evidencing, entails belief updating and belief updating incurs a thermodynamic cost via the Jarzynski equality [73, 74, 75]. This cost provides a lower bound on the thermodynamic free energy required for metabolic maintenance. For example, a cell's actions on its environment – e.g., chemotactic locomotion – are largely driven by the need to acquire thermodynamic free energy. The cell's GM cannot, therefore, minimize VFE by minimizing action [76]; instead, it must successfully predict which actions will replenish its free-energy supply. As actions are energetically expensive, this requires trading off short-term costs against long-term goals. As shown in [41], selective pressures operating on different timescales favor the development of metaprocessors that control lower-level actions in a context-dependent way; these are often implemented via a hierarchical GM [77]. Such meta-level control provides probabilistic models of risk-sensitive actions in context.

While such systems may be described as regulating free-energy seeking actions, they also regulate information-seeking actions, i.e., curiosity-driven exploration [78, 79, 80]. This follows because VFE provides an upper bound on complexity minus accuracy [81]. The expected free energy (EFE), conditioned upon any action, can therefore be scored in terms of expected complexity and expected inaccuracy. Expected complexity is "risk" and corresponds to the degree of belief updating that incurs a thermodynamic cost; leading to risk-sensitive control (e.g., phototropism). Expected inaccuracy corresponds to "ambiguity" leading to epistemic behaviors (e.g., searching for lost keys under a streetlamp) [42].

When context-dependent control is considered, the neighborhood of the NESS resolves into a network of local minima corresponding to fixed perception-action loops separated by energetic barriers that the control system must overcome to switch between loops. For example, in a cell, this energetic barrier comprises the energy required to activate one pathway while de-activating another, which may include the energetic costs of phosphorylation, other chemical modifications, additional gene expression, etc. Different pairs of pathways can be expected to be separated by energetic barriers of different heights, generating a topographically-complex free energy landscape that coarse-grains, in a long-time average, to the neighborhood of the NESS, i.e., to the maintenance of allostasis [68, 69, 82].

As noted earlier, we can think of controllable perception-action loops as nodes on a factor graph, with the edges corresponding to pathways for control flow, and the transition probabilities labeling the edges as inversely proportional to the energetic barrier between loops. This allows representing the GM for meta-level (i.e., hierarchical) control as a message-passing system as described in [47]. The presence of very high energetic barriers can render such a GM effectively one-way, as seen in the context-dependent switches between signal transduction pathways and GRNs that characterize cellular differentiation during morphogenesis. Biological examples of these include modifications of bioelectric pattern memories in planaria, which can create alternative-species head shapes that eventually remodel back to normal [83], or produce 2-headed worms which are permanent, and regenerate as 2-headed in perpetuity [84].

2.2 The QRF picture

Cellular information processing has traditionally been treated as completely classical, i.e., as implemented by causal networks of macromolecules, each of which undergoes classical state transitions via local dynamical processes that are conditionally independent of the states of other parts of the network. While the "quantum" nature of proteins and other macromolecules is broadly acknowledged, the scale at which quantum effects are important remains controversial, with straightforward single-molecule decoherence models predicting decoherence times of attoseconds (10^{-18} s) or less [85, 86]: several orders of magnitude below the timescales of processes involved in molecular information processing [87]. While functional roles for quantum coherence in intramolecular information processing have been demonstrated, intermolecular coherence remains experimentally elusive [88, 89, 90, 91].

The free-energy budgets of both prokaryotic and eukaryotic cells are, however, orders of magnitude smaller than would be required to support fully-classical information processing at the molecular scale, suggesting that cells employ quantum coherence as a computational resource [92]. Indirect evidence of longer-range, tissue-scale coherence in brains has also been reported [93]. Reformulating the FEP in quantum information-theoretic terms enables it to describe situations in which long-range coherence, and hence quantum computation, cannot be neglected.

Following the development in [12], we consider a bipartite decomposition $U = AB$ of a finite, isolated system U for which the interaction Hamiltonian $H_{AB} = H_U - (H_A + H_B)$ is sufficiently weak over the time period of interest that the joint state $|U\rangle$ is separable (i.e., factors) as $|U\rangle = |A\rangle|B\rangle$. In this case, we can choose orthogonal basis vectors $|i^k\rangle$ so that:

$$H_{AB} = \beta_k K_B T_k \sum_i^N \alpha_i^k M_i^k, \tag{2}$$

where K_B denotes Boltzmann's constant, T is the absolute temperature of the environment, $k =$ A or B, the M_i^k are N mutually-orthogonal Hermitian operators with eigenvalues in $\{-1, 1\}$, the $\alpha_i^k \in [0, 1]$ are such that $\sum_i^N \alpha_i^k = 1$, and $\beta_k \geq \ln 2$ is an inverse measure of k's thermodynamic efficiency that depends on the internal dynamics H_k; see [56, 58, 94, 95] for further motivation and details of this construction and [96] for a pedagogical review. This description is purely topological, attributing no geometry to either U or $\mathscr{B}$; hence it allows the "embedding space" of perceived "objects" to be an observer-dependent construct. It has several relevant consequences:

- We can regard A and B as separated, and determined by independent measures. They are separated by – and interact via – a holographic screen $\mathscr{B}$ that can be represented, without loss of generality, by an array of N non-interacting qubits, where N is the dimension of H_{AB} [94, 95].
- A and B can be regarded as exchanging finite N-bit strings, each of which encodes one eigenvalue of H_{AB} [94].
- A and B have free choice of basis for H_{AB}, corresponding to free choice of local frames at $\mathscr{B}$, e.g., free choice, for each qubit q_i on $\mathscr{B}$, of the local z axis and hence the z-spin operator s_z that acts on q_i [96].
- Choice of basis corresponds to choosing the zero-point of total energy by each of A and B. The systems A and B are, therefore, in general at informational, but not at thermal equilibrium [12].
- As A and B must obtain from B or A, respectively, whatever thermodynamic free energy is required, by Landauer's principle [73, 99, 100], to fund the encoding of classical bits on $\mathscr{B}$ (as well as any other irreversible classical computation), A and B must each devote some sector F of $\mathscr{B}$ to free-energy acquisition. The bits in F are "burned as fuel" and so do not contribute input data to computations. Waste-heat dissipation by one system is free energy acquisition by the other. The free-energy sectors F_A and F_B of A and B need not align as subsets of qubits on $\mathscr{B}$; that is,

qubits that A regards as free-energy sources may be regarded by B as informative outputs and vice-versa [56, 58].

- The actions of the internal dynamics H_A and H_B on $\mathscr{B}$ can be represented by A- and B-specific sets of QRFs, each of which both "measures" and "prepares" qubits on $\mathscr{B}$. Each QRF acts on the qubits in some specific sector of $\mathscr{B}$, breaking the permutation symmetry of Eq. (2) [56, 58, 59]. Only QRFs acting on sectors other than F implement informative computations; we will therefore restrict attention to these QRFs.

- Each "computational" QRF can, without loss of generality, be represented by a cone-cocone diagram (CCCD) comprising Barwise-Seligman classifiers and infomorphisms between them [54, 55]. The apex of each such CCCD is, by definition, both the category-theoretic limit and colimit of the "input/output" classifiers that correspond, formally, to the operators M_i^k in Eq. (2) [56, 58, 59].

Typically, a CCCD is structured as a distributed information flow in the form:

$$\begin{tikzcd}
\mathcal{A}_1 \arrow[r, leftrightarrow, "g_{21}", "g_{12}"'] & \mathcal{A}_2 \arrow[r, leftrightarrow, "g_{32}", "g_{23}"'] & \dots\ \mathcal{A}_k \\
& \mathbf{C}' \arrow[ul, "h_1"'] \arrow[u, "h_2"] \arrow[ur, "h_k"] & \\
\mathcal{A}_1 \arrow[ur, "f_1"] \arrow[r, leftrightarrow, "g_{21}", "g_{12}"'] & \mathcal{A}_2 \arrow[u, "f_2"'] \arrow[r, leftrightarrow, "g_{32}", "g_{23}"'] & \dots\ \mathcal{A}_k \arrow[ul, "f_k"']
\end{tikzcd} \tag{3}$$

incorporating sets of classifiers $\{A_\alpha\}$ and (logic) infomorphisms $\{f_i, g_{jk}\}$ [54, Ch 12] over suitable index ranges. As a memory-write system, Diagram (3) depicts a generic blueprint for a bow-tie or variational autoencoder (VAE) network amenable to describing a hierarchical Bayesian network with belief-updating as discussed in e.g. [12, 57, 59]. Crucially, it is the non-commutativity of CCCDs of this form that specifies intrinsic or quantum contextuality, as occurs, for instance, when the colimit core $\mathbf{C}'$ is undefinable [57, §7, §8] [59, §7.2]. Consequences of such contextuality are discussed via examples in Part II.

The holographic screen $\mathscr{B}$ functions as an MB separating A from B. It can be regarded as having an N-dimensional, N-qubit Hilbert space $\mathcal{H}_{q_i} = \prod_i q_i$. While $\mathcal{H}_{q_i}$ is strictly ancillary to $\mathcal{H}_U = \mathcal{H}_A \otimes \mathcal{H}_B$, the classical situation can be recovered in the limit in which the entanglement entropies $\mathcal{S}(|A\rangle), \mathcal{S}(|B\rangle) \to 0$ by considering the products $\mathcal{H}_A \otimes \mathcal{H}_{q_i}$ and $\mathcal{H}_B \otimes \mathcal{H}_{q_i}$ to be "particle" state spaces for A and B, respectively. In this classical limit, the states of $\mathcal{H}_{q_i}$ become the blanket states of an MB that functions as a classical information channel [94, 95, 96]. In quantum holographic coding, for example, $\mathscr{B}$ is often represented by a polygonal tessellation of the hyperbolic disc, with qubits represented by polygonal centroids. A specific TN model of a pentagon code is developed in [97]; see in particular their Fig. 4. The geometric description of $\mathscr{B}$ as implementing holographic coding, and its classical limit as an MB structured as a direct acyclic graph (DAG), is further explored in the setting of TQNNs in [98].

In this quantum-theoretic picture, "systems" or "objects" observed and manipulated by A or B correspond to sectors on $\mathscr{B}$ that are the domains of particular QRFs deployed by A or B, respectively [58, 12, 59]. To simplify notation, we use the same symbol, e.g., 'Q' to denote both a QRF Q and the sector dom(Q) on $\mathscr{B}$. Any identifiable system X factors into a "reference" component R that maintains a time-invariant state $|R\rangle$ or more generally, state density ρ_R, that allows re-identification and hence sequential measurements over extended time, and a "pointer" component P with a time-varying state $|P\rangle$ or density ρ_P. It is this pointer component, named for the pointer of an analog instrument, which is the "state of interest" for measurements. The QRFs R and P clearly must commute, and the sectors R and P clearly must be mutually decoherent [58, 12, 59]. All "system" sectors must be components of some overall sector E that corresponds to the "observable environment." The recording of measurement outcomes to a classical memory and the reading of previously-recorded outcomes from memory can similarly be represented by a QRF Y. As dom(Y) is a sector on $\mathscr{B}$, recorded memories of A are exposed to and hence subject to modification by B and vice-versa. Both the observable environment E and the memory sector Y must be disjoint from, and decoherent with, the free-energy sector F.

As actions on $\mathscr{B}$ encode classical data, they have an associated free energy cost of at least ln2 $K_B T$ per bit [73, 99, 100] that must originate from the source at F. Time-energy complementary associates a minimum time of $h/[\ln 2(K_B T)]$, with h being Planck's constant, to this energy expenditure. We can, therefore, associate actions on $\mathscr{B}$, including memory writes, with "ticks" of an internal time QRF, which we denote t_A and t_B for A and B, respectively. Assuming all observational outcomes are written to memory, we can represent the situation as in Fig. 1. The time QRF is effectively an outgoing bit counter that can be represented by a groupoid operator $\mathcal{G}_{ij} : t_i \rightarrow t_j$ [56]. As outgoing bits are oriented in opposite directions with respect to $\mathscr{B}$ for A and B, the time "arrows" t_A and t_B point in opposite directions. Hence A and B can both be regarded as "interacting with their own futures" as discussed in [96].

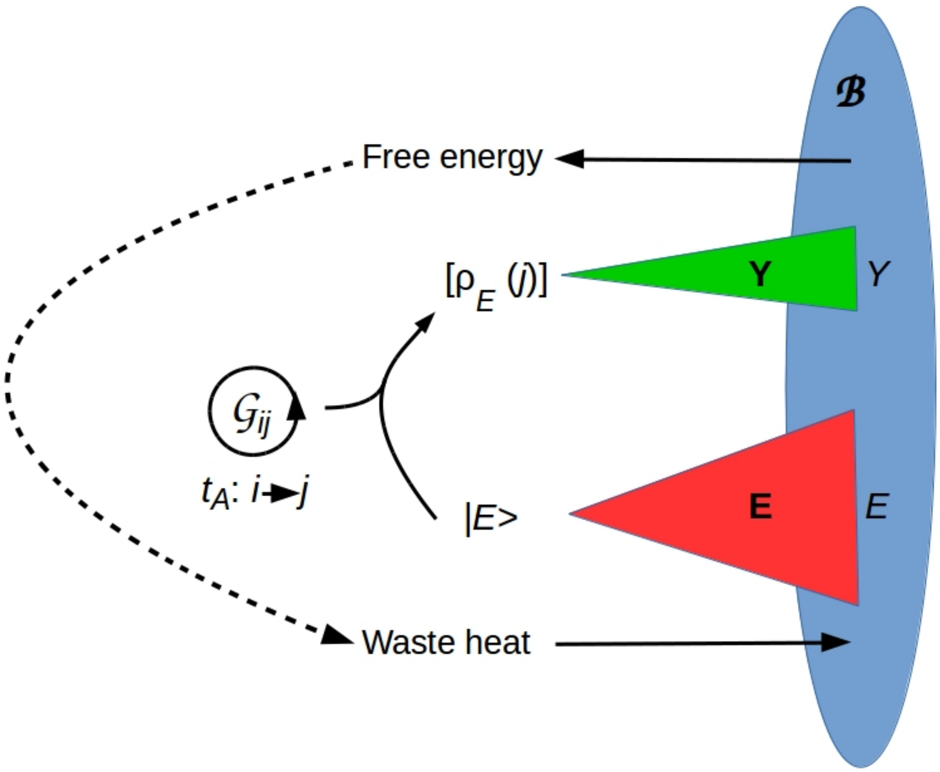

Figure 1: Cartoon illustration of QRFs required to observe and write a readable memory of an environmental state $|E\rangle$. The QRFs $\mathbf{E}$ and $\mathbf{Y}$ read the state from E and write it to the memory Y respectively. Any identified system S must be part of E. The clock $\mathcal{G}_{ij}$ is a time QRF that defines the time coordinate t_A. The dashed arrow indicates the observer's thermodynamic process that converts free energy obtained from the unobserved sector F of $\mathscr{B}$ to waste heat exhausted through F. Adapted from [58], CC-BY license.

Measurements of a system X can be considered sequential if: 1) they are separated in time according to the internal time QRF, and 2) their outcomes are recorded to memory to enable comparability across time. We show in [59] that sequential measurements can always be represented by one of two schemata. Using the compact notation:

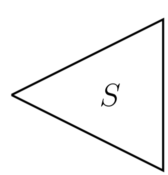

(4)

to represent a QRF S, we can represent measurements of a physical situation in which one system divides into two, possibly entangled, systems with a diagram of the form:

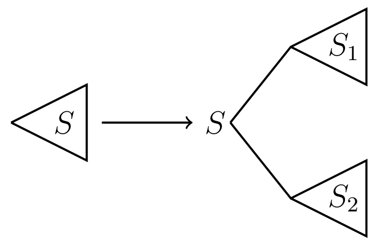

(5)

Parametric down-conversion of a photon exemplifies this kind of process. The reverse process can be added to yield:

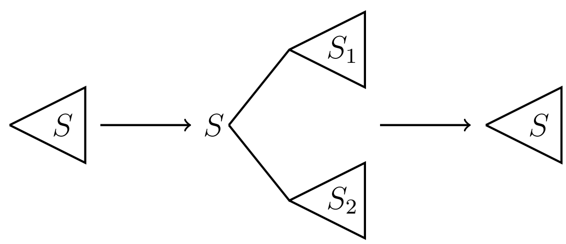

(6)

In the second type of sequential measurement process, the pointer-state QRF P is replaced with an alternative QRF Q with which it does not commute. Sequences in which position and momentum, or spins s_z and s_x, are measured alternately are examples. These can be represented by the diagram:

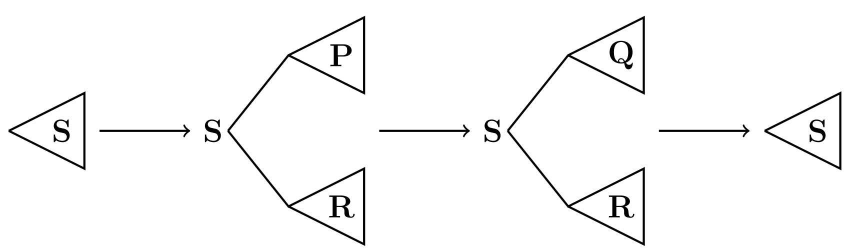

(7)

As both P and Q must commute with R, the commutativity requirements for S are satisfied.

The sequences of operations depicted in Diagrams (6) and (7) clearly raise the questions of how control is implemented, and of how the context changes that drive control flow are detected. Before turning to these questions in Part II, we review a path-integral representation of QRFs, show that the same representation also captures the behavior of any system X identified by a QRF, and discuss the questions of multiple observers and quantum contextuality.

2.3 The TQFT picture

As a least-action principle, the FEP is fundamentally a statement about the paths followed by the joint system U through its state space. The classical FEP is amenable to a path-integral formulation [13] that expresses the expected value of any observable (functional) $\Omega[x(t)]$ of paths $x(t)$ through the relevant state space as ([101], Eq. 6):

$$\langle \Omega[x(t)] \rangle = \int dx_0 \int d[x(t)] \Omega[x(t)] p(x(t)|x_0) p_0(x_0) \tag{8}$$

where x_0 is the initial state and $p(x(t)|x_0)$ is the conditional probability of the path $x(t)$. Quantum theory generalizes this expression by, effectively. replacing $\Omega[x(t)]$ with an automorphism on the relevant Hilbert space and $p(x(t)|x_0)$ with an amplitude for $x(t)$ given the initial state x_0. For some finite-dimensional Hilbert space $\mathcal{H}$, the manifold of all such automorphisms is a cobordism on $\mathcal{H}$, which is by definition a TQFT on $\mathcal{H}$ [102].

We show in [59] that any sequential measurement of any sector X of $\mathscr{B}$ induces a TQFT on X, considered as a projection of the N-dimensional boundary Hilbert space $\mathcal{H}_{q_i}$ associated with $\mathscr{B}$. In particular, measurement sequences of the form of Diagram (6) can be mapped to cobordisms, i.e., to manifolds of maps between two designated boundaries, of the form:

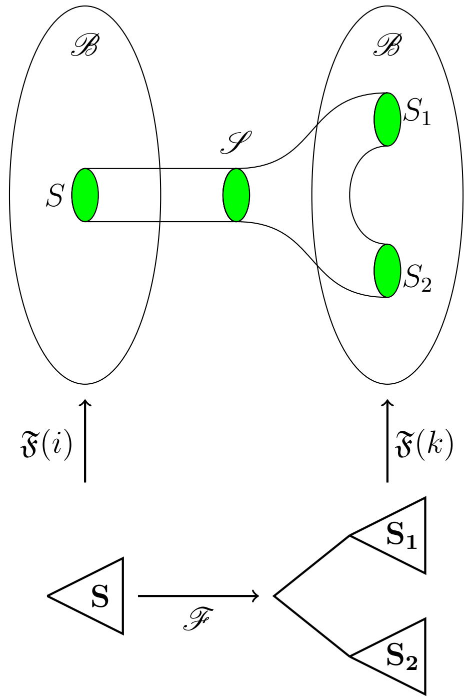

(9)

while sequences of the form of Diagram (7) can be mapped to cobordisms of the form:

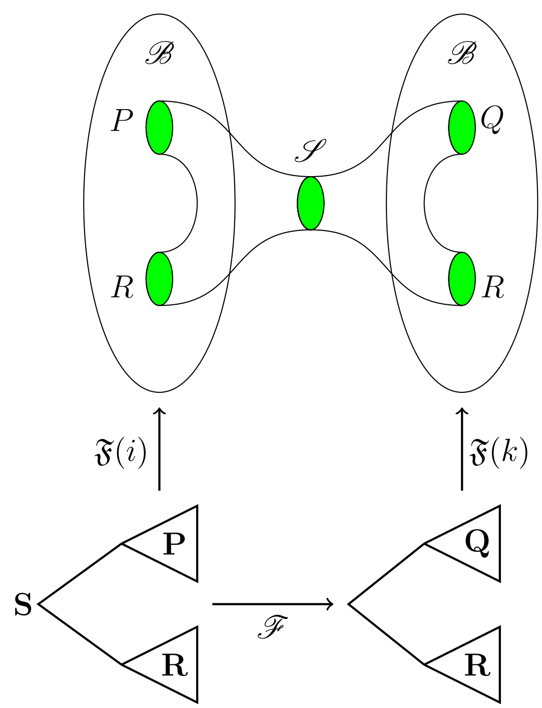

(10)

In either case, $\mathfrak{F} : \mathbf{CCCD} \to \mathbf{Cob}$ is the functor from the category **CCCD** of CCCDs (and hence of QRFs) to the category **Cob** of finite cobordisms required to define a TQFT. In general, we can state:

Theorem 1 ([59] Thm. 1)**.** *For any morphism $\mathscr{F}$ of CCCDs in* **CCCD***, there is a cobordism $\mathscr{S}$ such that a diagram of the form of Diagram* (9) *or* (10) *commutes.*

referring to [59] for the proof.

Theorem 1 applies to any sequential measurement; therefore, it applies to measurements of a sector X followed by measurements of the associated memory sector Y, or vice versa. Assuming for convenience that the dimension $\dim(X) = \dim(Y)$, we can consider a composite operation $Q = (\overrightarrow{Q}, \overleftarrow{Q})$, where $\overrightarrow{Q} = Q_X Q_Y$ and $\overleftarrow{Q} = Q_Y Q_X$. This Q is a pair of QRF sequences that can be identified with TQFTs that measure and record an outcome, mapping $\mathcal{H}_X \rightarrow \mathcal{H}_Y$, and dually use an outcome read from memory to prepare a state, mapping $\mathcal{H}_Y \rightarrow \mathcal{H}_X$, respectively, as in Diagram 11:

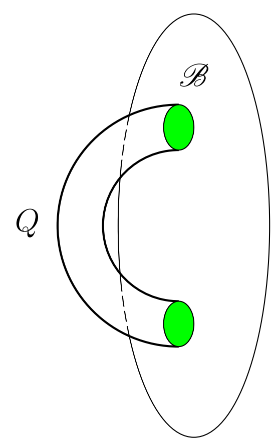

(11)

This composite operator Q is, by Theorem 1, itself a TQFT [98]. Hence the operation of recording observational outcomes for a sector X made at t to memory, and then comparing them to later observations at $t + \Delta t$, is formally equivalent to propagating the "system" X forward in time from t to $t + \Delta t$.

Identifying QRFs as "internal" TQFTs allows a general analysis of information exchange between multiple QRFs deployed by a single system, e.g., A. Because all QRFs act on $\mathscr{B}$, information exchange between QRFs requires a channel that traverses B. Any such channel is itself a QRF, one deployed by B. Considering A to comprise two observers, one deploying Q_1 and the other deploying Q_2, that interact via a local operations, classical communication (LOCC [103]) protocol provides an example:

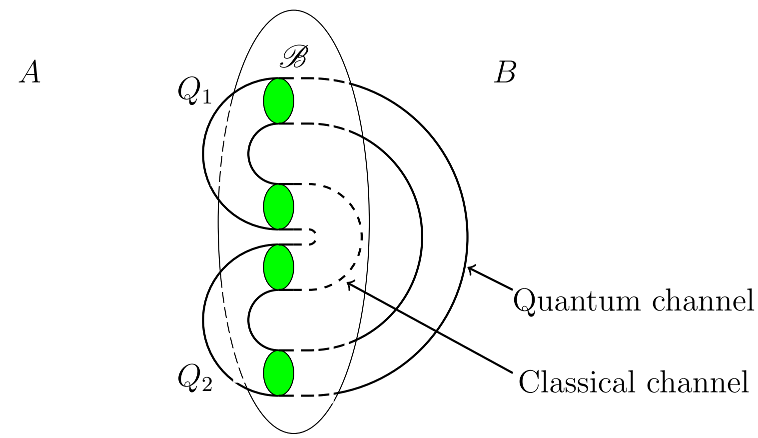

(12)

In a LOCC protocol, one channel is considered "classical" while the other is considered "quantum"; however, this language masks the fact that both channels are physical. As pointed out in [104], all media supporting classical communication are physical, and interactions with these media are always local measurements or preparations. Hence the two channels in a LOCC protocol are physically equivalent – both are TQFTs implemented by B – although their conventional semantics are different.

Diagram (12) can, clearly, also represent externally-mediated communication between any two functional components of a system, e.g., macromolecular pathways within a cell or functional networks within a brain. We show in [98] that whenever Q_1 and Q_2 are deployed by distinct – technically, separable or mutually decoherent – "observers" or "systems," they fail to commute, i.e., the commutator $[Q_1, Q_2] = Q_1Q_2 - Q_2Q_1 \geq h/2$, where again h is Planck's constant. As shown in [57], Theorem 3.4 using the CCCD representation, non-commutativity of QRFs induces quantum contextuality, i.e., dependence of measurement results on “non-local hidden variables” that characterize the measurement context [105, 106, 107]. In the current context, such hidden variables characterize the action of H_B on $\mathscr{B}$, affecting what A will observe next in every cycle of A-B interaction.

As shown in [63], such context dependence can, in principle, be captured classically if sufficient measurements of the context can be implemented. Such measurements would, however, have to access all of B. The existence of an MB prevents such access; in the current setting, A has access to B only via $\mathscr{B}$. The finite energetic cost of measurement, and consequent requirement for a thermodynamic sector F, prevents measurement even of all of $\mathscr{B}$ by any finite physical system. Hence, we can expect physical systems, including all biological systems, to employ only local context-dependent control to switch between mutually non-commuting (sets of) QRFs. How context switches implemented by QRF switches induce evolution, development and learning was introduced in [22]. Some specific examples of context switching in biological systems will be discussed Part II.

3 Conclusion

We have shown in this Part I how the problem of defining control flow arises in active inference systems, and provided three formal representations of the problem. Control flow can, in particular, be represented as switching between classical dynamical attractors, between deployed QRFs, and between computational processes represented by TQFTs. Implementing control flow has a free-energy cost; hence any control-flow system must trade off its own processing costs against the expected benefits of switching between input/ouput modes. The time and memory dependence of control flow can, moreover, be expected to lead generically to context effects on both perception and action.

In the accompanying Part II of this paper, we will first prove that control flows in active inference systems can always be represented as TNs, and show how TN architectures provide a convenient classification control flows. We then show how these can be implemented by TQNNs, and discuss applications of this formalism to the problem of characterizing control flow in biological systems.

Acknowledgements

K.F. is supported by funding for the Wellcome Centre for Human Neuroimaging (Ref: 205103/Z/16/Z), a Canada-UK Artificial Intelligence Initiative (Ref: 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 John Templeton Foundation, Grant 62230. 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.

References

[1] Friston KJ. A theory of cortical responses. *Philos Trans R Soc Lond B, Biol Sci* 2005;360:815–36.

[2] Friston KJ, Kilner J, Harrison L. A free energy principle for the brain. *J Physiol Paris* 2006;100:70–87.

[3] Friston KJ, Stephan KE. Free-energy and the brain. *Synthese* 2007;159:417–58.

[4] Friston, K. J. 2010 The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience 11*, 127–138.

[5] Friston, K. J. 2013 Life as we know it. *Journal of The Royal Society Interface 10*, 20130475.

[6] Friston KJ, FitzGerald T, Rigoli F, Schwartenbeck P, Pezzulo G. Active inference: a process theory. *Neural Comput* 2017;29:1–49.

[7] Ramstead MJ, Badcock PB, Friston KJ. Answering Schrödinger's question: a free-energy formulation. *Phys Life Rev* 2018;24:1–16.

[8] Ramstead MJ, Constant A, Badcock PB, Friston KJ. 2019 Variational ecology and the physics of sentient systems. *Phys Life Rev* 31, 188–205.

[9] Kuchling F, Friston K, Georgiev G, 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.

[10] Friston, K. J. 2019 A free energy principle for a particular physics. Preprint arxiv:1906.10184 [q-bio.NC]. https://arxiv.org/abs/1906.10184

[11] Ramstead MJ, Sakthivadivel DAR, Heins C, Koudahl M, Millidge B, Da Costa L, Klein B, Friston KJ 2022 On Bayesian mechanics: A physics of and by beliefs. *Interface Focus* 13, 2022.0029.

[12] Fields C, Friston K, Glazebrook JF, Levin M 2022 A free energy principle for generic quantum systems. *Prog. Biophys. Mol. Biol.* 173, 36–59.

[13] Friston, K., Da Costa, L., Sakthivadivel, D. A. R., Heins, C., Pavliotis, G. A., Ramstead, M., Parr, T. 2022 Path integrals, particular kinds, and strange things. Preprint arxiv:2210.12761.

[14] Pearl, J. 1988 *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* San Mateo CA: Morgan Kaufmann.

[15] Clark A. 2017 How to knit your own Markov blanket: Resisting the second law with metamorphic minds. In (T. Wetzinger and W. Wiese, eds.) *Philosophy and Predictive Processing 3*, 19pp. Frankfurt am Main: Mind Group.

[16] Kirchhoff, M., Parr, T., Palacios, E., Friston, K., Kiverstein, J. 2018 The Markov blankets of life: Autonomy, active inference and the free energy principle. *J. R. Soc. Interface* 15, 20170792.

[17] Parr, T., Da Costa, L., Friston, K. 2020 Markov blankets, information geometry and stochastic thermodynamics. *Philos. Trans. A: Math. Phys. Eng. Sci.* 378, 20190159.

[18] Sakthivadivel, D.A.R. 2022 Weak Markov blankets in high-dimensional, sparsely-coupled random dynamical systems. Preprint arXiv:2207.07620.

[19] Wheeler, J. H. 1989 Information, physics, quantum: The search for links. In: Zurek, W. (Ed.), *Complexity, Entropy, and the Physics of Information.* CRC Press, Boca Raton, FL, pp. 3–28.

[20] Levin, M. 2019 The computational boundary of a “self”: Developmental bioelectricity drives multicellularity and scale-free cognition. *Front. Psychol. 10*, 1688.

[21] 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*, 114–133.

[22] Fields, C.; Glazebrook, J. F.; Levin, M. 2021 Minimal physicalism as a scale-free substrate for cognition and consciousness. *Neurosci. Cons. 7(2)*, niab013.

[23] Levin, M. 2022 Technological approach to mind everywhere: An experimentally-grounded framework for understanding diverse bodies and minds. *Front. Syst. Neurosci.* 16, 768201.

[24] Endsley, M. R. 2012 Situational awareness. In: Salvendy, G. (Ed.) *Handbook of Human Factors and Ergonomics*, 4th Ed. Hoboken, NJ, John Wiley, pp. 553–568.

[25] Fields, C.; Levin, M. 2018 Multiscale memory and bioelectric error correction in the cytoplasm-cytoskeleton-membrane system. *WIRES Syst. Biol. Med.* 10, e1410.

[26] Clark, A., Chalmers, D. 1998 The extended mind (Active externalism). *Analysis* 58(1), 7–19.

[27] Anderson, M. L. 2003 Embodied cognition: A field guide. *Artif. Intell.* 149, 91–130.

[28] Froese, T., Ziemke, T 2009 Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artif. Intell.* 173, 466–500.

[29] Attias, H. 2003 Planning by probabilistic inference. *Proc. of the 9th Int. Workshop on Artificial Intelligence and Statistics* in Proc. Machine Learning Res. R4, 9–16.

[30] Botvinick, M., Toussaint, M. 2012 Planning as inference. *Trends Cogn. Sci.* 16(10), 485–488.

[31] Lanillos, P., Mio, C., Pezzato, C., et al. 2021 Active inference in robotics and artificial agents: Survey and challenges. Preprint arXiv:2112.01871.

[32] Chubukov, V., Gerosa, L., Kochanowski, K., Sauer, U. 2014 Coordination of microbial metabolism. *Nat. Rev. Microbiol.* 12, 327–340.

[33] Micali, G., Endres, R. G. 2016 Bacterial chemotaxis: information processing, thermodynamics, and behavior. *Curr. Opin. Microbiol.* 30, 8–15.

[34] Pezzulo, G., LaPalme, J., Durant, F., Levin, M. 2021 Bistability of somatic pattern memories: stochastic outcomes in bioelectric circuits underlying regeneration. *Philos. Trans. R. Soc. Lond. B* 376(1821), 20190765.

[35] Blake, R., Logothetis, N. K. 2002 Visual competition. *Nat. Rev. Neurosci.* 3, 1–11.

[36] Stertzer, P., Kleinschmidt, A., Rees, G. 2009 The neural bases of multistable perception. *Trends Cogn. Sci.* 13, 310–318.

[37] Schwartz, J.-L., Grimault, N., Hupé, J.-M., Moore, B. C. J., Pressnitzer, D. 2012 Multistability in perception: bindingsensory modalities, An overview. *Phil. Trans. R. Soc. Lond. B* 367, 896–905.

[38] Vossel, S., Geng, J. J., Fink, G. R. 2014 Dorsal and ventral attention systems: Distinct neural circuits but collaborative roles. *Neuroscientist* 20, 150–159.

[39] Baars, B. J., Franklin, S. 2003 How conscious experience and working memory interact. *Trends Cogn. Sci.* 7, 166–172.

[40] Baars, B. J., Franklin, S., Ramsoy, T. Z. 2013 Global workspace dynamics: Cortical "binding and propagation" enables conscious contents. *Front. Psychol.* 4, 200.

[41] Kuchling, F.; Fields, C.; Levin, M. 2022 Metacognition as a consequence of competing evolutionary time scales. *Entropy* 24, 601.

[42] Parr, T., Friston, K. J. 2019 Generalised free energy and active inference. *Biol. Cybern.* 113(5-6), 495–513.

[43] Winn, J. Bishop, C. M. 2005 Variational message passing. *J. Mach. Learn. Res.* 6, 661–694.

[44] Dauwels, J. 2007 On variational message passing on factor graphs. *2007 IEEE International Symposium on Information Theory*, Nice, France.

[45] Parr, T., Sajid, N., Friston, K. J. 2020 Modules or mean-fields? *Entropy* 22(5), 552.

[46] Da Costa, L., Parr, T. Sajid, N., Veselic, S., Neacsu, V., Friston, K. 2020 Active inference on discrete state-spaces: A synthesis. *J. Math. Psychol.* 99, 102447.

[47] Friston, K., Parr, T., de Vries, B. 2017 The graphical brain: Belief propagation and active inference. *Netw. Neurosci.* 1(4), 381–414.

[48] Orús, R. 2019 Tensor networks for complex quantum systems. *Nat. Rev. Phys.* 1, 538–550.

[49] Bao, N., Cao, C.-J., Carroll, S. M., Chatwin-Davies, A. 2017 de Sitter space as a tensor network: Cosmic no-hair, complementarity, and complexity. *Phys. Rev. D* 96, 123536.

[50] Hu, Q., Vidal, G. 2017 Spacetime symmetries and conformal data in the continuous Multiscale Entanglement Renormalization Ansatz. *Phys. Rev. Lett.* 119, 010603.

[51] Chandra, A. R., de Boer, J., Flory, M., Heller, M. P., Hörtner, S., Rolph, A. 2021 Spacetime as a quantum circuit. *J. High Energy Phys.* 2021, 207.

[52] Aharonov Y, Kaufherr T. Quantum frames of reference. *Phys Rev D* 1984;30:368–385.

[53] Bartlett SD, Rudolph T, Spekkens RW. 2007 Reference frames, super-selection rules, and quantum information. *Rev Mod Phys* 79, 555–609.

[54] Barwise, J.; Seligman, J. 1997 *Information Flow: The Logic of Distributed Systems* (Cambridge Tracts in Theoretical Computer Science 44). Cambridge University Press, Cambridge, UK.

[55] Fields, C.; Glazebrook, J. F. 2019 A mosaic of Chu spaces and Channel Theory I: Category-theoretic concepts and tools. *J. Expt. Theor. Artif. intell. 31*, 177–213.

[56] Fields, C.; Glazebrook, J. F. 2020 Representing measurement as a thermodynamic symmetry breaking. *Symmetry 12*, 810.

[57] Fields, C.; Glazebrook, J. F. 2022 Information flow in context-dependent hierarchical Bayesian inference. *J. Expt. Theor. Artif. intell. 34*, 111–142.

[58] Fields, C.; Glazebrook, J. F.; Marcianò, A. 2021 Reference frame induced symmetry breaking on holographic screens. *Symmetry 13*, 408.

[59] Fields, C.; Glazebrook, J. F.; Marcianò, A. (2022) Sequential measurements, topological quantum field theories, and topological quantum neural networks. *Fortschr. Phys.* 2022, 2200104.

[60] Abramsky, S., Brandenburger, A. 2011 The sheaf-theoretic structure of non-locality and contextuality. *New J. Phys.* 13, 113036.

[61] Abramsky, S., Barbosa, R. S., Mansfield, S. 2017 Contextual fraction as a measure of contextuality. *Phys. Rev. Lett.* 119, 050504.

[62] Dzhafarov, E. N.; Kujala, J. V. (2017a). Contextuality-by-Default 2.0: Systems with binary random variables. In: J. A. Barros, B. Coecke and E. Pothos (eds.) *Lecture Notes in Computer Science* 10106, Springer, Berlin, 16–32.

[63] Dzharfarov, E. N. & Kon, M. 2018 On universality of classical probability with contextually labeled random varaibles. *J. Math. Psychol.* 85, 17–24.

[64] Adlam, E. 2021 Contextuality, fine-tuning and teleological explanation. *Found. Phys.* 51, 106.

[65] Marcianò, A.; Chen, D.; Fabrocini, F.; Fields, C.; Greco, E.; Gresnigt, N.; Jinklub, K.; Lulli, M., Terzidis, K.; Zappala, E. 2022 Quantum neural networks and topological quantum field theories. *Neural Networks* 153, 164–178.

[66] Marcianò, A., Chen, D., Fabrocini, F., Fields, C., Lulli, M., Zappala, E. 2022 Deep neural networks as the semi-classical Limit of topological quantum neural networks: The problem of generalisation. Preprint arXiv:2210.13741.

[67] Sterling, P., Eyer, J. 1988 Allostasis: A new paradigm to explain arousal pathology. In: *Handbook of Life Stress, Cognition and Health.* John Wiley & Sons: New York, pp. 629–649.

[68] Barrett, L. F., Quigley, K. S., Hamilton, P. 2016 An active inference theory of allostasis and interoception in depression. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 371(1708), 20160011.

[69] Corcoran, A. W., Pezzulo, G., Hohwy, J. 2020 From allostatic agents to counterfactual cognisers: Active inference, biological regulation, and the origins of cognition. *Biol. Philos.* 35(3), 32.

[70] Hohwy, J. 2016 The self-evidencing brain. *Noûs* 50(2), 259–285.

[71] Sakthivadivel, D. A. R. 2022 A constraint geometry for inference and integration. Preprint arXiv:2203.08119.

[72] Sakthivadivel, D. A. R. 2022 Towards a geometry and analysis for Bayesian mechanics. Preprint arXiv:2204.11900.

[73] Landauer, R. 1961 Irreversibility and heat generation in the computing process. *IBM J. Res. Dev. 5*, 183–195.

[74] Jarzynski, C. 1997 Nonequilibrium equality for free energy differences. *Phys. Rev. Lett.* 78(14), 2690–2693.

[75] Evans, D. J. 2003 A non-equilibrium free energy theorem for deterministic systems. *Molec. Phys* 101(10), 1551–1554.

[76] Friston, K., Thornton, C., Clark, A. 2012 Free-energy minimization and the dark-room problem. *Front. Psychol.* 3, 130.

[77] Pezzulo, G., Rigoli, F., Friston, K. 2015 Active Inference, homeostatic regulation and adaptive behavioural control. *Prog. Neurobiol.* 134, 17–35.

[78] Friston, K., Rigole, F., Ognibene, D., Mathys, C., Fitzgerald, T., Pezzulo, G. 2015 Active inference and epistemic value. *Cogn. Neurosci.* 6, 187-214.

[79] Schmidhuber, J. 1991 Curious model-building control-systems. *1991 IEEE International Joint Conference on Neural Networks*, Vols 1-3; 2, 1458–1463.

[80] Sun, Y., Gomez, F., Schmidhuber, J. 2011 Planning to be surprised: Optimal Bayesian exploration in dynamic environments. *Artificial General Intelligence.* J. Schmidhuber, K. R. Thórisson and M. Looks, Eds. Berlin, Heidelberg, Springer. pp. 41–51.

[81] Sengupta, B., Friston, K. 2018 How Robust are deep neural networks? Preprint arXiv arXiv:1804.11313.

[82] Seth, A. K., Friston, K. J. 2016 Active interoceptive inference and the emotional brain. *Philos. Trans. R. Soc. Lond. B* 371(1708), 20160007.

[83] Emmons-Bell, M., Durant, F., Hammelman, J. et al. 2015 Gap junctional blockade stochastically induces different species-specific head anatomies in genetically wild-type *Girardia dorotocephala* flatworms. *Int. J. Mol. Sci.* 16, 27865–27896.

[84] Oviedo, N. J., J. Morokuma, Walentek, P. et al. 2010 Long-range neural and gap junction protein-mediated cues control polarity during planarian regeneration. *Dev. Biol.* 339, 188–199.

[85] Tegmark M 2000 Importance of quantum decoherence in brain processes. *Phys. Rev. E* 61: 4194–4206.

[86] Schlosshauer M. 2007 *Decohenece and the Quantum to Classical Transition.* Springer, Berlin.

[87] Zweir MC, Chong LT. 2010 Reaching biological timescales with all-atom molecular dynamics simulations. *Curr. Opin. Pharmacol.* 10: 745–752.

[88] Marais A *et al.* 2018 The future of quantum biology. *J. R. Soc. Interface* 15: 20180640.

[89] Cao J *et al.* 2020 Quantum biology revisited. *Science Adv.* 6: eaaz4888.

[90] Kim, Y., Bertagna, F., D'Souza, E. M., Heyes, D. J., Johannissen, L. O., Nery, E. T. 2021 Quantum biology: an update and perspective. *Quant. Rep.* 3, 1–48.

[91] Baiardi, A., Christandl, M., Reiher, M. 2022 Quantum computing for molecular biology. Preprint arxiv:2212.12220.

[92] Fields, C.; Levin, M. 2021 Metabolic limits on classical information processing by biological cells. *BioSystems 209*, 104513.

[93] Kerskens, C. M., Pérez, D. L. 2022 Experimental indications of non-classical brain functions. *J. Phys. Commun.* 6, 105001.

[94] Fields, C.; Marcianò, A. 2019 Holographic screens are classical information channels. *Quant. Rep.* 2, 326–336.

[95] Addazi, A.; Chen, P.; Fabrocini, F.; Fields, C.; Greco, E.; Lulli, M.; Marcianò, A.; Pasechnik, R. 2021 Generalized holographic principle, gauge invariance and the emergence of gravity à la Wilczek. *Front. Astron. Space Sci.* 8, 563450.

[96] Fields, C.; Glazebrook, J. F.; Marcianò, A. 2022 The physical meaning of the holographic principle. *Quanta* 11, 72–96.

[97] Pastawski, F.; Yoshida, B.; Harlow, D.; Preskill, J. 2015 Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence. *J. High Energy Phys.* 6, 149.

[98] Fields, C.; Glazebrook, J. F.; Marcianò, A. 2023 Communication protocols and quantum error-correcting codes from the perspective of topological quantum field theory. Preprint arxiv:2303.16461 [hep-th].

[99] Landauer, R. 1999 Information is a physical entity. *Physica A* 263, 63–67.

[100] Bennett, C. H. 1982 The thermodynamics of computation. *Int. J. Theor. Phys.* 21, 905–940.

[101] Seifert, U. 2012 Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.* 75, 126001.

[102] Atiyah, M. 1988 Topological quantum field theory. *Pub. Math. IHÈS* 68, 175–186.

[103] Chitambar, E., Leung, D., Mančinska, L., Ozols, M., Winter, A. 2014 Everything you always wanted to know about LOCC (but were afraid to ask). *Comms. Math. Phys.* 328, 303–326.

[104] Tipler, F. 2014 Quantum nonlocality does not exist. *Proc. Natl. Acad. Sci. USA* 111, 11281–11286.

[105] Bell, J. S. 1966 On the problem of hidden variables in quantum mechanics. *Rev. Mod. Phys.* 38, 447–452.

[106] Kochen, S., Specker, E. P. 1967 The problem of hidden variables in quantum mechanics. *J. Math. Mech.* 17, 59–87.

[107] Mermin, N. D. 1993 Hidden variables and the two theorems of John Bell. *Rev. Mod. Phys.* 65, 803–815.