

Article

A Variational Synthesis of Evolutionary and Developmental Dynamics

Karl Friston ¹, Daniel A. Friedman ^{2,3,*}, Axel Constant ⁴, V. Bleu Knight ^{3,5}, Chris Fields ⁶, Thomas Parr ¹ and John O. Campbell ⁷

¹ Wellcome Centre for Human Neuroimaging, Institute of Neurology, University College London, London WC1E 6AP, UK; k.friston@ucl.ac.uk (K.F.)

² Department of Entomology and Nematology, University of California, Davis, CA 95616, USA

³ Active Inference Institute, Davis, CA 95616, USA; blanket@activeinference.institute

⁴ Theory and Method in Biosciences, The University of Sydney, Sydney, NSW 2006, Australia; axel.constant.pruvost@gmail.com

⁵ Department of Biology, New Mexico State University, Las Cruces, NM 88003, USA

⁶ Allen Discovery Center at Tufts University, Medford, MA 02155, USA; fieldsres@gmail.com

⁷ Independent Researcher, Victoria, BC V8V 1X4, Canada

* Correspondence: danielariefriedman@gmail.com

Abstract: This paper introduces a variational formulation of natural selection, paying special attention to the nature of ‘things’ and the way that different ‘kinds’ of ‘things’ are individuated from—and influence—each other. We use the Bayesian mechanics of particular partitions to understand how slow phylogenetic processes constrain—and are constrained by—fast, phenotypic processes. The main result is a formulation of adaptive fitness as a path integral of phenotypic fitness. Paths of least action, at the phenotypic and phylogenetic scales, can then be read as inference and learning processes, respectively. In this view, a phenotype actively infers the state of its econiche under a generative model, whose parameters are learned via natural (Bayesian model) selection. The ensuing variational synthesis features some unexpected aspects. Perhaps the most notable is that it is not possible to describe or model a population of conspecifics *per se*. Rather, it is necessary to consider populations of distinct natural kinds that influence each other. This paper is limited to a description of the mathematical apparatus and accompanying ideas. Subsequent work will use these methods for simulations and numerical analyses—and identify points of contact with related mathematical formulations of evolution.

Keywords: self-organisation; nonequilibrium; variational inference; Bayesian; particular partition; evolution; natural selection; Markov blanket; renormalisation group

Citation: Friston, K.; Friedman, D.A.; Constant, A.; Knight, V.B.; Fields, C.; Parr, T.; Campbell, J.O. A Variational Synthesis of Evolutionary and Developmental Dynamics. *Entropy* **2023**, *25*, 964. <https://doi.org/10.3390/e25070964>

Academic Editors: Damián H. Zanette and Rainer Klages

Received: 20 March 2023

Revised: 12 June 2023

Accepted: 15 June 2023

Published: 21 June 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Dedicated to the Memory of John O. Campbell.

1. Introduction

This paper is an attempt to show that some fundamentals of theoretical evolution—and (neuro)biology—emerge when applying the free energy principle to dynamical systems with separation of temporal scales. It offers a technical and generic treatment with minimal assumptions or commitments to specific biological processes. As such, it does not borrow from established constructs in evolutionary theory; rather, it tries to show how some of these constructs are emergent properties, when seen through the lens of the free energy principle. In subsequent work, we will use the ensuing variational synthesis to consider established—and current—evolutionary theories. Our aim in this paper is to introduce a formalism that may be useful for addressing specific questions—about evolutionary or developmental dynamics—using analytic or numerical recipes that have proven useful when applying the free energy principle in other fields.

A key phylogenetic process—underlying the development and diversification of species in evolutionary time—is known as natural selection, regarded by some as the central organizing principle of biology. While Darwin conceived of natural selection in terms of heredity, variation, and selection [1,2], he only detailed selection, as the mechanisms of heredity and variation would not be understood for some time [3,4]. The integration of Mendelian genetics with natural selection in the early twentieth century was followed by an integration with molecular genetics [5] in the mid-century to form Neo-Darwinism, or the modern synthesis. The modern synthesis, along with the selfish gene hypothesis—put forth in the 1970s [6]—provide a largely gene-centric view of Darwinian evolution that dominates the current perspective.

This gene-centric view of evolutionary biology has remained largely disconnected from phenotypic processes that impact organisms in developmental time [7,8]. Lewontin characterised this disconnect—between genetic and phenotypic understanding—as the major challenge facing the field [9]. While some progress has been made in the following fifty years, biologists continue to highlight the gaps remaining for modelling biology as a single integrated process over multiple scales [10–13]. By ‘gene-centric’, we refer not just to theories of sequence evolution [14], but also to the central role genes (or summary statistics of genes) play either explicitly or implicitly in accounts of phenotypic evolution. For instance, the Price Equation [15] and the closely related replicator equation [16] of evolutionary game theory express the relationship between the changes in (the average of) some phenotypic trait over time. This gene-centric view relies upon a mapping between that trait and the genetic material passed from generation to generation but focuses upon the phenotypic effects of genes as opposed to the alleles themselves. Similarly, adaptive dynamic approaches [17] typically focus upon ecological interactions at a phenotypic level. The modern focus upon phenotypic traits reflects the importance of the interaction between a phenotype and its environment in determining fitness. However, it is important to note that such perspectives do not conflict with the central role of genetic inheritance, and implicitly score the fitness of genotypes in terms of the phenotypes they imply.

An organism inherits a set of instructions for growth and development (i.e., an extended genotype) that is, in essence, a prediction about the niche environment (including temperature, humidity, chemical composition, available resources, statistical patterns, etc.). Interrogating the phrase ‘survival of the fittest’ leads to the understanding of ‘fittest’ as organisms that are the best ‘fit’ to their niche environment [18]. For example, a bacterium from thermal hot springs will fail to thrive in a cool pond because its genotype does not accurately predict the niche environment. Therefore, ‘fitness’ is relative to the niche, where slow phylogenetic processes have selected for an extended genotype that enhances the growth and proliferation of organisms in the environment where the corresponding species expects to find itself.

An organism can also ‘fit’ itself to the niche through adaptation (i.e., action, learning, and development) during its lifetime. For example, a bacterium that normally subsists on sulphur reduction—but can also survive through reducing oxygen—will outlast its sulphur-dependent competitors in an environment that is devoid of sulphur. Such an organism can adapt to its environment through learning and optimising for oxygen reduction, thereby increasing its fit to the niche and, implicitly, its capacity to reproduce in a high-oxygen environment. In this way, the phenotypic processes can enhance the fit of organisms to their environment in developmental time, and through reproduction, phenotypic processes can lead to the enhancement of fit in evolutionary time (i.e., across generations). As the (extended) genotype of organisms produces phenotypes, phylogenetic processes over evolutionary time also impact phenotypic (ontogenetic) processes in developmental time.

Here, we offer a synthesis of evolution and development through a mathematical framework that unifies slow, multi-generational (phylogenetic) processes with single-life-time, phenotypic (developmental and behavioural) processes using the same principles, as they apply to each temporal scale. The ensuing variational account of evolution focuses

on the coupling between phylogenetic processes at evolutionary timescales and ontogenetic processes over phenotypic lifetimes. In principle, this abstract treatment is agnostic to specific mechanisms, and could be applied to biological as well as non-biological systems provided their ‘fitness’ depends upon events during a lifetime, and where this fitness influences dynamics over a generational scale. This multiscale account foregrounds the circular causality that arises from the implicit separation of timescales [19].

In brief, we consider slow phylogenetic processes (*natural selection*) as furnishing top-down constraints (i.e., top-down causation) on fast phenotypic processes (*action selection*). In turn, the active exchange of the phenotype with its environment provides evidence that is assimilated by natural selection (i.e., bottom-up causation). This ontological account is licensed by describing both phylogenetic and phenotypic processes as selecting (extended) genotypes and (extended) phenotypes [7,20] with the greatest fitness, where fitness is quantified with (free energy) functionals of probability density functions (a functional is a function of a function).

This formulation means that natural selection and action selection can be described as updating probabilistic beliefs at phylogenetic and phenotypic scales, respectively: namely, learning and inference [21–23]. This separation of scales affords an interpretation of *natural selection as Bayesian model selection* [24–26], while *action selection becomes planning as inference* [27–30]—both (appearing to) optimise the same fitness functional: namely, Bayesian model evidence or marginal likelihood. A narrative version of this account can be told from the point of view of the genotype (from the bottom up) or the phenotype (from the top down):

From the perspective of the genotype, we can consider evolution as belief-updating over generations, where the *belief* in question corresponds to a probability density over extended genotypes (henceforth, genotype). This belief-based model of allelic change is analogous to treatments of evolution in terms of changes in allele frequencies from generation to generation [15]. This belief updating can be described by the probability of a genotype appearing in subsequent generations, in a way that depends lawfully on the marginal likelihood of extended phenotypes (henceforth, phenotype) in the current generation. *The basic idea is that the genotype parameterises or encodes a generative model, which the phenotype uses to infer and act on its environment.* On this view, evolution can be regarded as testing hypotheses—in the form of generative models—that this kind of phenotype can persist in this environment. These hypotheses are tested by exposing the phenotype to the environment and are rejected if the phenotype ‘strays from the path’ of a persistent phenotype. In this way, the evolutionary process selects models or hypotheses about persistent phenotypes for which it has the greatest evidence. In short, natural selection is just Bayesian model selection [25,26,31,32].

From the perspective of a phenotype, each conspecific is equipped with a generative model and initial conditions that underwrite its epigenetic, developmental and ethological trajectories. The states of the phenotype trace out a path through state-space over its lifetime. These phenotypic states encode or parameterise beliefs about environmental states—and the way the phenotype acts. This parameterization leads to active inference and learning, in which the phenotype tries to make sense of its world and—through a process of belief updating—to realise the kind of creature it thinks it is. (We use the term ‘thinks’ in a liberal sense here and do not mean to imply that all living entities have explicit existential thoughts.) More precisely, what we mean is that these entities behave as if they hold a set of beliefs about the sort of entity they are (e.g., the meta-Bayesian stance as considered in [33]). In virtue of its genetic endowment, it thinks it is a persistent phenotype. If endowed with a good generative model of its environment [34], it will persist and supply evidence of its ‘fit’ to the environment (i.e., ‘fitness’); namely, evidence (i.e., marginal likelihood) that has been accumulated by the slow evolutionary process.

What follows is a formal version of this narrative that calls upon some standard results from statistical physics. The resulting synthesis is both dense and delicate, because it tries to account for coupling between a phenotype and its econiche—and the coupling

between phenotypic and phylogenetic processes—using the same principles. Specifically, we describe a variational synthesis that calls on the path integral formulation of stochastic dynamics, the apparatus of the renormalisation group, and the Poincaré recurrence theorem. The ensuing synthesis considers natural selection and action selection as emergent properties of two random dynamical processes unfolding at slow (phylogenetic) and fast (phenotypic) timescales. The key aspect of this synthesis is that both processes have an attracting set (a.k.a., pullback attractor) or steady-state solution [35]. These solutions correspond to an *evolutionary stable state* [36] and a *nonequilibrium steady-state density* [37] over phylogenetic and phenotypic states, respectively. By describing these steady states in terms of a phylogenetically encoded generative model—namely, a joint density over the paths of the phenotype and its environment—one can recover an ontological description of how the two processes inform, and are informed by, each other.

Some of the analysis presented in this paper follows that in [21–23], which also appeals to the notion of a renormalisation group. These treatments are based upon the emergence of separable timescales and the interpretation of the dynamics at each scale in analogy with inference and learning processes. The key differences are as follows. The renormalisation in [21] depends upon a reduction in the number of degrees of freedom with learning, whereas our formulation depends upon a partitioning operation as part of the renormalisation. The difference in timescales between variables in [21] emerges from the structure of the neural network used, whereas it is a direct consequence of the reduction operator implicit in our choice of renormalisation. Finally, we extend our analysis to sentient phenotypes, whose dynamics can be interpreted explicitly in terms of Bayesian belief-updating. We conclude with a numerical study, illustrating the basic ideas with synaptic selection in the brain.

2. A Variational Formulation

We assume that evolution can be described with two random dynamical systems, describing phylogenetic (evolutionary) and phenotypic (particular) processes, respectively. The idea is to couple these systems using the apparatus of the renormalisation group [38–40] to map from fast phenotypic dynamics to slow phylogenetic dynamics in evolutionary time.

This mapping rests upon a dimension reduction and coarse graining or grouping operator (**RG** for Renormalisation Group) that maps the path of a phenotype $\tilde{\pi}$ to relevant variables at the evolutionary scale $\bar{\pi} = \mathbf{R} \circ \tilde{\pi}$. On this view, bottom-up causation is simply the application of a reduction operator, $\mathbf{R} \circ \tilde{\pi}$, to select variables that change very slowly. Top-down causation entails a specification of fast phenotypic trajectories in terms of slow genotypic variations, which are grouped into populations, $\mathbf{G} \circ \bar{\pi}$, according to the influences they exert on each other. The implicit separation into fast and slow variables can be read as an adiabatic approximation [41] or—in the sense of synergetics—into fast (dynamically stable) and slow (dynamically unstable) modes, respectively [42]. This separation can also be seen in terms of vectorial geometric formulations [43]. Please see [21], who deal carefully with the separation of time scales by analogy with temporal dilation in physics. Intuitively, this analogy rests upon the idea that time can be rescaled, depending upon whether we take the perspective of things that move quickly or slowly.

The final move is to express the dynamics—at fast and slow levels—in terms of functionals that have the same form. These functionals are functions of probability densities that can be read as Bayesian beliefs. Expressing the dynamics in this way allows one to interpret phenotypic dynamics as active inference and learning, under a generative model that depends on the extended genotype. In other words, one can interpret the phylogenetic state as inferring states of the environment over evolutionary time. Crucially, the extended genotype accumulates evidence for its phenotype, thereby evincing a form of Bayesian model selection or structure learning [25,44–48]. For an analogous thermodynamic treatment, please see [22], who refine and extend the free energy formulation of

[49]. In the context of learning dynamics, a thermodynamic free energy was derived in [50]—using the maximum entropy principle [51,52]—and later applied to study phenomenological models of evolution [22]. Please see [50,53,54] for further discussion in terms of neural networks and information theory.

2.1. Particular Partitions

There are many moving parts in this formulation because it tries to account for the behaviour of ‘things’ [55] and how this behaviour underwrites the emergence of ‘kinds’ (e.g., individuals and populations) at nested (i.e., developmental and evolutionary) time-scales.

We will use $[x(t)] \subset \tilde{x}$ to denote the history or path of a time-varying state. These paths are determined by state-dependent flow $f_{\tilde{x}}(x)$, with parameters $\bar{x} \subset \tilde{x}$ that include initial states $x(0) = x_0 \subset \bar{x}$. These parameters denote a (natural) kind.

Everything that follows rests upon a *particular partition* of states. A particular partition is considered necessary to talk about ‘things’, such as a ‘phenotype’ or ‘population’. In brief, a particular partition enables the (internal) states of some ‘thing’ to be separated from the (external) states of every ‘thing’ else by (sensory and active) blanket states [56–60]. In the absence of this partition, there would be no way of distinguishing a phenotype from its external milieu—or a population from the environment. In this setup, external states can only influence themselves and sensory states, while internal states can only influence themselves and active states. See Figure 1 for an influence diagram representing the coupling among internal, external, and blanket states:

States: $x = (\eta, s, a, \mu)$. States comprise the *external*, *sensory*, *active* and *internal* states of a phenotype. Sensory and active states constitute *blanket* states $b = (s, a)$, while *phenotypic* states comprise internal and blanket states, $\pi = (b, \mu) = (s, \alpha)$. The *autonomous* states of a phenotype $\alpha = (a, \mu)$ are not influenced by external states:

- i. **External states** respond to sensory and active states. These are the states of a phenotype’s external milieu: e.g., econiche, body, or extracellular space, depending upon the scale of analysis.
- ii. **Sensory states** respond to fluctuations in external and active states: e.g., chemo-reception, proprioception, interception.
- iii. **Active states** respond to sensory and internal states and mediate action on the environment, either directly or vicariously through sensory states: e.g., actin filaments, motor action, autonomic reflexes.
- iv. **Internal states** respond to sensory and active states: e.g., transcription, intracellular concentrations, synaptic activity.




Figure 1. Schematic (i.e., influence diagram) illustrating the sparse coupling among states that constitute a particular partition at two scales.

The evolution of these sparsely coupled states can be expressed as a Langevin or stochastic differential equation: namely, a high dimensional, nonlinear, state-dependent flow plus independent random (Wiener) fluctuations, ω , with a variance of 2Γ :

$$\dot{x} = f_{\bar{x}}(x) + \omega = \begin{bmatrix} \dot{\eta} \\ \dot{s} \\ \dot{a} \\ \dot{\mu} \end{bmatrix} = \begin{bmatrix} f_{\bar{\eta}}(\eta, s, a) \\ f_{\bar{s}}(\eta, s, a) \\ f_{\bar{a}}(s, a, \mu) \\ f_{\bar{\mu}}(s, a, \mu) \end{bmatrix} + \begin{bmatrix} \omega_{\eta} \\ \omega_s \\ \omega_a \\ \omega_{\mu} \end{bmatrix} \quad (1)$$

The flow per se can be expressed using the Helmholtz–Hodge decomposition [61] as follows:

$$f_{\bar{x}}(x) = (Q - \Gamma) \nabla \mathfrak{I}(x)$$

$$\begin{bmatrix} f_{\bar{\eta}}(\eta, b) \\ f_{\bar{s}}(\eta, b) \\ f_{\bar{a}}(b, \mu) \\ f_{\bar{\mu}}(b, \mu) \end{bmatrix} = \begin{bmatrix} Q_{\eta} - \Gamma_{\eta} & Q_{\eta s} \\ -Q_{\eta s}^T & Q_s - \Gamma_s \\ & & Q_a - \Gamma_a & Q_{a\mu} \\ & & -Q_{a\mu}^T & Q_{\mu} - \Gamma_{\mu} \end{bmatrix} \begin{bmatrix} \nabla_{\eta} \mathfrak{I}(\eta | b) \\ \nabla_s \mathfrak{I}(b | \eta) \\ \nabla_a \mathfrak{I}(b | \mu) \\ \nabla_{\mu} \mathfrak{I}(\mu | b) \end{bmatrix} \quad (2)$$

Note that our appeal to an equation of this form means we have implicitly stipulated that there is a steady-state density or potential function that remains constant (or at least changes very slowly) over the timescale we are interested in. Equation (2) expresses the flow as a mixture of a dissipative, gradient flow and a conservative, solenoidal flow [62–64]. The gradient flow $\Gamma \nabla \mathfrak{I}$ depends upon the amplitude of random fluctuations, while

the solenoidal flow $Q\nabla\mathfrak{I}$ circulates on the isocontours of the potential function called *self-information*, $\mathfrak{I}(x) = -\ln p(x)$, where $p(x)$ is called the *nonequilibrium steady-state density* or NESS density [37,65–67].

The particular partition above rests on sparse coupling between dynamic variables, c.f., [68,69], and evinces the notion of an ‘action-perception cycle’ between external and internal states [70]. The terms ‘external’ and ‘internal’ offer useful intuitions, but it is worth being cautious about overinterpreting these labels in spatial terms. For instance, it might seem that some ‘external’ variables such as ambient temperature might directly influence ‘internal’ variables such as the temperature within a cell. However, this intuition would not be an appropriate way of thinking about this system’s partition. Either we would have to assume that there is an intervening variable (e.g., the temperature within the cell membrane) or we would have to treat the internal temperature as a sensory variable, which itself influences internal variables such as the rates of enzymatic reactions. There is now an emerging literature asking about the appropriate ways to think of particular partitions in biology, including what is internal to a neuronal network [71], or a spinal reflex arc [72].

2.2. Ensemble Dynamics and Paths of Least Action

To describe dynamics at the phenotypic or phylogenetic scale, we first need to rehearse some standard results from statistical physics that furnish a probabilistic description of trajectories or paths at any scale. This description calls on the self-information of states $x(t)$, generalised states $\vec{x} = (x, x', \dots)$, and paths, $\tilde{x} = [x(t)]$, where $\mathbf{D}\vec{x} = (x', x'', \dots)$ denotes generalised notion, and $2\mathbf{\Gamma}$ is the covariance of generalised random fluctuations:

$$\begin{aligned}\mathfrak{I}(x) &= -\ln p(x) \\ \mathcal{L}(\vec{x}) &= -\ln p(\vec{x} | x_0) = \frac{1}{2}[\ln |\mathbf{\Gamma}| + (\mathbf{D}\vec{x} - f(\vec{x})) \cdot \frac{1}{2\mathbf{\Gamma}} (\mathbf{D}\vec{x} - f(\vec{x})) + \nabla \cdot f] \\ \mathcal{A}(\tilde{x}) &= -\ln p(\tilde{x} | x_0) = \int dt \mathcal{L}(\vec{x})\end{aligned}\quad (3)$$

The first measure, $\mathfrak{I}(x)$, is the self-information or *surprisal* of a state, namely, the implausibility of a state being occupied. When the state is an allele frequency and evolves according to Wright–Fisher dynamics, this is sometimes referred to as an ‘adaptive landscape’ [73]. The second, $\mathcal{L}(\vec{x})$, is the *Lagrangian*, which is the *surprisal* of a generalised state, namely, the instantaneous path associated with the motion from an initial state. In generalised coordinates of motion, the state, velocity, acceleration, etc., are treated as separate (generalised) states that are coupled through the flow [74,75]. Finally, the surprisal of a path $\mathcal{A}(\tilde{x})$ is called *action*, namely, the path integral of the Lagrangian.

Generalised states afford a convenient way of expressing the path of least action in terms of the Lagrangian

$$\nabla_{\vec{x}} \mathcal{L}(\vec{x}) + (\dot{\vec{x}} - \mathbf{D}\vec{x}) = 0 \Leftrightarrow \dot{\vec{x}} - \mathbf{D}\vec{x} = -\nabla_{\vec{x}} \mathcal{L}(\vec{x}) \Leftrightarrow \dot{\vec{x}}(\tau) = \mathbf{D}\vec{x} - \nabla_{\vec{x}} \mathcal{L}(\vec{x}) \quad (4)$$

The first equality resembles a Lagrange equation of the first kind that ensures the generalised motion of states is the state of generalised motion. Alternatively, it can be read as a gradient descent on the Lagrangian, in a moving frame of reference (second equality). When the Lagrangian is convex, solutions to this generalised gradient descent on the Lagrangian (third equality) necessarily converge to the path of least action. Denoting paths of least action with boldface:

$$\begin{aligned}\dot{\vec{x}} &= \mathbf{D}\vec{x} = f(\vec{x}) \\ \Leftrightarrow \nabla_{\vec{x}} \mathcal{L}(\vec{x}) &= 0 \Leftrightarrow \vec{x} = \arg \min_{\vec{x}} \mathcal{L}(\vec{x}) \\ \Leftrightarrow \delta_{\vec{x}} \mathcal{A}(\vec{x}) &= 0 \Leftrightarrow \vec{x} = \arg \min_{\vec{x}} \mathcal{A}(\vec{x})\end{aligned}\quad (5)$$

Convergence is guaranteed by the quadratic form (i.e., convexity) of the Lagrangian, which inherits from Gaussian assumptions about random fluctuations. This gradient descent is sometimes described as convergence to the path of least action, in a frame of reference that moves with the state of generalised motion [76].

We can also express the conditional independencies implied by a particular partition using the Lagrangian of generalised states. Because there are no flows that depend on both internal and external states, external and internal paths are independent, when conditioned on blanket paths:

$$\frac{\partial^2 f}{\partial \eta \partial \mu} = 0 \Rightarrow \frac{\partial^2 \mathcal{L}}{\partial \vec{\eta} \partial \vec{\mu}} = 0 \Leftrightarrow \mathcal{L}(\vec{\eta}, \vec{\mu} | \vec{s}, \vec{a}) = \mathcal{L}(\vec{\eta} | \vec{s}, \vec{a}) + \mathcal{L}(\vec{\mu} | \vec{s}, \vec{a}) \Leftrightarrow (\vec{\eta} \perp \vec{\mu}) | \vec{s}, \vec{a}, x_0 \quad (6)$$

In other words, blanket paths furnish a Markov blanket over internal paths. We will use this result later to disambiguate the role of active and sensory dynamics in sentient behaviour—i.e., active inference—of a phenotype. First, we have to establish a formalism for ensembles or populations of phenotypes. Here, we draw on the apparatus of the renormalisation group.

2.3. Different Kinds of Things

To deal with multiple ‘things’ (e.g., particles, phenotypes and populations), we first introduce a grouping operator $\mathbf{G}$ that partitions the states at the i -th scale of analysis into N particles on the basis of the sparse coupling implied by a particular partition. In other words, we group states into an ensemble of particles, where each particle has its own internal and blanket states. With a slight abuse of the set builder notation:

$$\{\pi_1^{(i)}, \dots, \pi_N^{(i)}\} = \mathbf{G} \circ \underbrace{\{x_1^{(i)}, \dots, x_j^{(i)}, \dots, x_k^{(i)}, \underbrace{x_\ell^{(i)}, \dots, x_m^{(i)}}_{a_n^{(i)}}, \underbrace{x_o^{(i)}, \dots, x_p^{(i)}}_{\mu_n^{(i)}}\}}_{\pi_n^{(i)}} \quad (7)$$

The grouping operator means the external states of a given particle are the (blanket) states of remaining particles that influence it. See [55] for a worked example and numerical analysis. This grouping expresses the dynamics of each particle in terms of its sensory states—that depend upon the blanket states of other particles—and autonomous states—that only depend upon the states of the particle in question:

$$\dot{\pi}_n = \begin{bmatrix} \dot{s}_n \\ \dot{\alpha}_n \end{bmatrix} = \begin{bmatrix} f_{\bar{s}_n}(b_n, \dots, b_N) \\ f_{\bar{\alpha}_n}(b_n, \mu_n) \end{bmatrix} + \omega_n \quad (8)$$

$$\begin{bmatrix} f_{\bar{s}_n} \\ f_{\bar{\alpha}_n} \end{bmatrix} = \begin{bmatrix} Q_{s_n} - \Gamma_{s_n} & 0 \\ 0 & Q_{\alpha_n} - \Gamma_{\alpha_n} \end{bmatrix} \begin{bmatrix} \nabla_{s_n} \mathfrak{Z}(b_1, \dots, b_N) \\ \nabla_{\alpha_n} \mathfrak{Z}(\alpha_n | s_n) \end{bmatrix}$$

At this point, we pause to consider that the states in the particular ensemble have to be the states of some ‘thing’: namely, the states of a particle at a lower scale. This means that states must be the states of particles (e.g., phenotypic states) that constitute the particular states at the next scale (e.g., phylogenetic states). This recursive truism can be

expressed in terms of grouping $\mathbf{G}$ operator—that creates particles—and a reduction $\mathbf{R}$ operator—that picks out certain particular states for the next scale:

$$\xrightarrow{\mathbf{R}} \{\tilde{x}_\ell^{(i)}\} \xrightarrow{\mathbf{G}} \{\tilde{\pi}_n^{(i)}\} \xrightarrow{\mathbf{R}} \{\tilde{x}_n^{(i+1)}\} \xrightarrow{\mathbf{G}} \{\tilde{\pi}_m^{(i+1)}\} \xrightarrow{\mathbf{R}} \quad (9)$$

The composition of the two operators can be read as mapping from the states of particles at one scale to the next or, equivalently, from particular states at one scale to the next—in short, creating particles of particles, namely, populations. See Figure 2.

$$\begin{aligned} & \xrightarrow{\mathbf{R} \circ \mathbf{G}} \{\tilde{x}_\ell^{(i)}\} \xrightarrow{\mathbf{R} \circ \mathbf{G}} \{\tilde{x}_n^{(i+1)}\} \xrightarrow{\mathbf{R} \circ \mathbf{G}} \\ & \xrightarrow{\mathbf{G} \circ \mathbf{R}} \{\tilde{\pi}_n^{(i)}\} \xrightarrow{\mathbf{G} \circ \mathbf{R}} \{\tilde{\pi}_m^{(i+1)}\} \xrightarrow{\mathbf{G} \circ \mathbf{R}} \end{aligned}$$

$$\{\pi(\tau)_1^{(i+1)}, \dots, \pi(\tau)_M^{(i+1)}\} = \mathbf{G} \circ \underbrace{\{\mathbf{R} \circ \tilde{\pi}_1^{(i)}, \dots, \mathbf{R} \circ \tilde{\pi}_j^{(i)}, \dots, \mathbf{R} \circ \tilde{\pi}_k^{(i)}, \mathbf{R} \circ \tilde{\pi}_\ell^{(i)}, \dots, \mathbf{R} \circ \tilde{\pi}_m^{(i)}, \dots\}}_{\substack{s_m^{(i+1)} \\ \pi_m^{(i+1)}}} \quad (10)$$

$$\begin{aligned} \tilde{\pi}_n^{(i)} &= \mathbf{R} \circ \tilde{\pi}_n^{(i)} \\ \{\pi_m^{(i+1)}\} &= \mathbf{G} \circ \{\tilde{\pi}_n^{(i)}\} \end{aligned}$$




Figure 2. Schematic showing the hierarchical relationship between particles at scales i and $i + 1$. For clarity, sensory and autonomous states are illustrated in blue and pink, respectively. Note that each variable is a (very large) vector state that itself is partitioned into multiple vector states. At scale $i + 1$, each particle represents an ensemble (e.g., $\pi_m^{(i+1)}$ is population m), the elements of which are partitioned into autonomous and sensory subsets (e.g., $\alpha_{m_n}^{(i+1)}$ is the n -th autonomous genotype from population m). At scale i , each particle represents an element of an ensemble (e.g., $\pi_\ell^{(i)}$ is the ℓ -th

phenotype), which is itself partitioned into sensory and autonomous subsets. The slow states of each element (e.g., phenotype) are recovered by the reduction operator $\mathbf{R}$, to furnish the states at the ensemble level (e.g., genotype). A key feature of this construction is that it applies recursively over scales.

The reduction operator $\mathbf{R}$ typically selects relevant variables whose slow fluctuations contextualise dynamics at the scale below. Here, $\mathbf{R}$ simply recovers the states of a particle that are time invariant or that vary slowly with time (i.e., the initial states and flow parameters). This separation of timescales means that the lifetime of a particle (e.g., phenotype) unfolds during an instant from the perspective of the next scale (e.g., evolution). The separation of timescales could have been achieved without the grouping (partitioning) operator. We could simply have projected onto the eigenvectors of a dynamical system's Jacobian, effectively taking linear (or nonlinear) mixtures of our system to arrive at fast and slow coordinates. However, all we would be left with are fast and slow continuous variables that have nothing of the character of the individuals, phenotypes, or populations in a system. In short, the grouping operator is key in identifying fast and slow 'things'—as opposed to just fast and slow coordinates of a dynamical system.

In short, the renormalisation group operator creates particles of particles, retaining only particular variables that change very slowly and then grouping them according to their sparse coupling. This means that particles increase in their size from one scale to the next—in virtue of the grouping of particles at the lower scale—and change more slowly—in virtue of the coarse graining afforded by temporal reduction.

In an evolutionary setting, the existence of steady-state solutions—implicit in the Langevin formalism above—means that phenotypic dynamics possess a pullback attractor. This means their paths will return to the neighbourhood of previously occupied states. In other words, their 'lifecycle' will intersect with some Poincaré section in phenotypic state-space (possibly many times). We will take this intersection to be a mathematical image of persistence, which is underwritten by the flow parameters at any point in evolutionary time.

At the phylogenetic scale, we have a partition into populations of phenotypes based upon which phenotypes influence each other. At this slow scale, states can be read as characterising the 'kind' of 'thing' that has particular states at the scale below. We will, therefore, refer to states at this level as (natural) kinds, noting that the 'kind of thing' in question does not change at the fast scale. We can now rehearse the particular partition at the phylogenetic scale, noting that for a population to exist, it must have a particular partition. Here, a population corresponds to a set of particular kinds $x^{(i+1)} = (\eta, s, a, \mu)$. These include *external*, *sensory*, *active*, and *internal* kinds.

- i. **External kinds** of particles are phenotypes outside the population that change as a function of themselves and sensory and active kinds: c.f., the target of *niche construction*, from a molecular through to a cultural level, depending upon the scale of analysis [77,78].
- ii. **Sensory kinds** mediate the effects of external kinds on the internal members of the population in question: e.g., nutrients or prey.
- iii. **Active kinds** mediate the effects of internal kinds on external kinds: e.g., agents who mediate niche construction, from a molecular through to a cultural level, depending upon the scale of analysis.
- iv. **Internal kinds** influence themselves and respond to changes in sensory and active kinds.

This concludes our formal setup. Next, we consider the coupling between fast phenotypic and slow phylogenetic dynamics. As in other applications of the free energy principle, this coupling emerges as a property of any phylogenetics that possesses an evolutionary steady state. In other words, the idea here is to identify the properties of a system

that exists, as opposed to identifying the properties that underwrite existence. We will see that the emergent properties look very much like natural selection.

2.4. Natural Selection: A Variational Formulation

To specialise particular partitions to natural selection, we will associate autonomous (active and internal) kinds with the (extended) genotypes that constitute a population of *agents*, noting that there is no requirement for agents to belong to the same equivalence class—they just interact, in virtue of the sparse coupling that defines their grouping into a population. For example, some agents could be animals, and others could be plants.

At the phylogenetic scale, an agent is an autonomous kind from a particular population. At the phenotypic scale, the agent has particular (phenotypic) states, whose dynamics or paths depend upon its (genotypic) kind. For ease of notation, we will deal with a single population where the phenotypic state of the n -th agent, $\alpha_n^{(i+1)}$, will be denoted by $\pi_\ell^{(i)}$ (i.e., dropping the m in Figure 2). With this formalism in place, we can formulate the coupling between phenotypic and phylogenetic dynamics with the following lemma:

Lemma 1. (Variational fitness): *If, at non-equilibrium evolutionary steady state, the likelihood of an agent's genotype $\alpha_n^{(i+1)} = \mathbf{R} \circ \tilde{\pi}_\ell^{(i)}$ is proportional to the likelihood of its phenotypic trajectory $\tilde{\pi}_\ell^{(i)}$ (where $\setminus$ denotes exclusion),*

$$\begin{aligned}\mathfrak{I}_n^{(i+1)} &= \mathfrak{I}(\alpha_n^{(i+1)} | \pi^{(i+1)} \setminus \alpha_n^{(i+1)}) \\ &= \mathcal{A}(\tilde{\pi}_\ell^{(i)} | \bar{x}^{(i)} \subset \pi^{(i+1)}) = \mathcal{A}_\ell^{(i)}\end{aligned}\quad (11)$$

then the following holds:

An agent's autonomous dynamics can be cast as a gradient descent on a Lagrangian, whose path integral (i.e., action) corresponds to negative fitness. This Lagrangian depends upon the flow parameters (and initial states) supplied by the genotype. The agent's genotype can then be cast as a stochastic gradient descent on negative fitness. This formulation emphasises the relationship between gradients on fitness (selection) and the stochastic terms that are uncorrelated with selection (drift):

$$\begin{aligned}\overbrace{\dot{\tilde{\alpha}}_\ell^{(i)} = \mathbf{D} \tilde{\alpha}_\ell^{(i)} - \nabla_{\tilde{\alpha}_\ell^{(i)}} \mathcal{L}_\ell^{(i)}}^{\text{Fast (c.f., phenotypic) dynamics}} &\iff \overbrace{\dot{\alpha}_n^{(i+1)} = (Q_{\alpha_n}^{(i+1)} - \Gamma_{\alpha_n}^{(i+1)}) \nabla_{\alpha_n^{(i+1)}} \mathcal{A}_\ell^{(i)} + \omega_n^{(i+1)}}^{\text{Slow (c.f., phylogenetic) dynamics}} \\ \underbrace{\mathcal{L}_\ell^{(i)} = \mathcal{L}(\tilde{\pi}_\ell^{(i)} | \bar{x}^{(i)} \subset \pi^{(i+1)})}_{\text{Lagrangian (c.f., surprisal)}} &\quad \quad \quad \underbrace{\mathcal{A}_\ell^{(i)} = \int dt \mathcal{L}(\tau)_\ell^{(i)}}_{\text{Action (c.f., adaptive fitness)}}\end{aligned}\quad (12)$$

Formally, the generalised gradient descent at the phenotypic scale corresponds to Bayesian filtering or *inference* [76] that maximises the marginal likelihood of phenotypic paths. This is almost tautological, in that it says that deviations from the most likely developmental trajectory, given some genotype, are unlikely. An additional subtlety here is that the Lagrangian, which plays the role of a Lyapunov function, is a function of sensory states. The implication is that the gradients are not static, but themselves change based upon the way in which the environment interacts with a creature during its development. The stochastic gradient descent at the phylogenetic scale corresponds to Bayesian *learning* via stochastic gradient Langevin dynamics [79], equipped with solenoidal mixing [80].

On this Bayesian reading, phenotypic dynamics *infer* their external dynamics, under a probabilistic model of how external dynamics generate phenotypic dynamics. Intergenerational genetic changes can be seen as *learning* the parameters of a generative model, given the Bayesian model evidence supplied by the scale below (e.g., extended phenotype). This reading rests upon the action (i.e., negative fitness) scoring the accumulated evidence $p(\tilde{\pi} | \bar{x})$ for a phenotype's generative model, $p(\tilde{\eta}, \tilde{\pi} | \bar{x})$ encoded by the

extended genotype $\bar{x}$. This evidence is also known as a *marginal likelihood* because it marginalises over external dynamics; i.e., other agents.

Proof. The condition in (11) means that the probability of finding an agent of a particular kind is proportional to the likelihood of its phenotypic path, namely, the likelihood a phenotype keeps to the ‘trodden path’, characteristic of the ‘kind’ of ‘things’ that persist. The existence of a nonequilibrium evolutionary steady-state solution to the density dynamics (at both scales) allows us to express the fast and slow dynamics of agents and their autonomous states in terms of Helmholtz–Hodge decompositions. From (1) and (2), we have

$$\begin{aligned} \dot{\alpha}_\ell^{(i)} &= (Q_\alpha^{(i)} - \Gamma_\alpha^{(i)}) \nabla_{\alpha_\ell^{(i)}} \mathfrak{S}_\ell^{(i)} + \omega_\ell^{(i)} \\ \mathfrak{S}_\ell^{(i)} &= \mathfrak{S}(\pi_\ell^{(i)} | \bar{x}_\ell^{(i)} \subset \pi^{(i+1)}) \end{aligned} \quad \Longleftrightarrow \quad \begin{aligned} \dot{\alpha}^{(i+1)} &= (Q_\alpha^{(i+1)} - \Gamma_\alpha^{(i+1)}) \nabla_{\alpha^{(i+1)}} \mathfrak{S}^{(i+1)} + \omega^{(i+1)} \\ \mathfrak{S}^{(i+1)} &= \mathfrak{S}(\pi^{(i+1)}) \end{aligned} \quad (13)$$

The gradients of surprisal at the slow scale, with respect to any given agent’s ‘kind’ or genotype, are the gradients of action by (11):

$$\nabla_{\alpha_n^{(i+1)}} \mathfrak{S}^{(i+1)} = \nabla_{\alpha_n^{(i+1)}} \mathfrak{S}_n^{(i+1)} = \nabla_{\alpha_n^{(i+1)}} \mathcal{A}_\ell^{(i)} \quad (14)$$

Substituting (14) into (13) gives the slow, phylogenetic dynamics in (12) (ignoring certain solenoidal terms). $\square$

For the fast, phenotypic dynamics, we assume that random fluctuations vanish to describe phenotypes that possess classical (i.e., Lagrangian) mechanics, i.e., that are dominated by conservative or solenoidal dynamics. In the limit of small fluctuations, the autonomous paths become the paths of least action, i.e., when the fluctuations take their most likely value of zero. From (4), the autonomous paths of least action are as follows (setting $\lambda = 1$):

$$\dot{\vec{\alpha}}_\ell^{(i)} = \mathbf{D} \vec{\alpha}_\ell^{(i)} - \nabla_{\vec{\alpha}_\ell^{(i)}} \mathcal{L}(\vec{\pi}_\ell^{(i)} | \bar{x}_\ell^{(i)}) \quad (15)$$

Substituting (15) into (13) gives the fast dynamics in (12).

Remark 1. Note that the extended genotype $\bar{x}_\ell^{(i)} = \{\bar{\eta}_\ell^{(i)}, \bar{\pi}_\ell^{(i)}\} \subset \pi^{(i+1)}$ includes the initial states of the extended phenotype. In other words, the extended genotype covers both the genetic and epigenetic specification of developmental trajectories and the initial conditions necessary to realise those trajectories, including external states (e.g., conditions necessary for embryogenesis), $\eta(0)_\ell^{(i)} \subset \bar{\eta}_\ell^{(i)}$.

A useful intuition as to the biological role of the Lagrangian in Equation (11) is that it specifies the states (or trajectories) of a system that has achieved homeostasis. The function will return a small value when physiological measurements are within homeostatic ranges, and increasingly large values as deviations from these ranges become larger. The conditioning upon slow (genotypic) variables means that different sorts of homeostatic ranges are allowable for different sorts of phenotypes. The relationship between the (fast) action and (slow) Lagrangian in Equation (11) implies that phenotypic trajectories—in which homeostasis is maintained—are associated with genotypes that are more likely to be replicated. More precisely, the Lagrangian favours (i.e., its path integral is smaller for) those trajectories in which opportunities for replication are attained—and successful maintenance of homeostasis is only one aspect of this.

The suppression of random phenotypic fluctuations does not preclude itinerant trajectories. Indeed, it foregrounds the loss of detailed balance and accompanying nonequilibria that characterise phenotypic and population dynamics [81–83]: for example, bio-rhythms and chaotic oscillations at the phenotypic scale [84–88] or Red Queen dynamics

at the phylogenetic scale [83,89,90]. A system that has the property of detailed balance is one in which time reversal makes no qualitative difference to the dynamics of that system. The implication is that systems in which the solenoidal flow is zero possess detailed balance, while those with a non-zero solenoidal flow do not. The presence of solenoidal flow means that time reversal also leads to a reversal in the direction of this flow. Please see [31] as a relatively recent example of the Helmholtz–Hodge decomposition in Darwinian processes and [80] for a generic treatment of stochastic chaos in this setting. Furthermore, there is no requirement for the grouping operator to return the same partition at each instant of its application. This follows because the grouping operator is determined by sparse coupling among particles at the scale below, which itself may change as certain particles become ‘shielded’ from others [91]: for example, during the self-assembly of particular partitions associated with cell-division, multicellular organisation and development [57]. Mathematically, this permits wandering sets (i.e., partitions) at each scale, where fitness gradients remain well-defined, because they inherit from the dynamics of the scale below.

Implicit in the renormalisation group construction is the notion that variational selection could operate at multiple scales. In other words, although framed in terms of natural selection and evolution, the variational formulation above does not commit to separation of temporal scales apt for replication or reproduction. Any selective mechanism that fulfils the fitness lemma (Lemma 1) will, in principle, be subject to the same selective mechanics. Common examples could include the optimisation of weights in neural networks and their structure learning [45,76,92]. In a biological setting, this selection process could correspond to developmental stages that have well-defined (separation of) temporal scales. Finally, we take a closer look at phenotypic dynamics and explain why they can be construed as sentient behaviour.

3. The Sentient Phenotype

An ontological interpretation of phenotypic dynamics—in terms of sentient behaviour or active inference—obtains by expressing the Lagrangian as a *variational free energy*. For clarity, we will drop the sub- and superscripts (and condition on the extended genotype $\bar{x}$) to focus on the generalised states of a given phenotype.

Lemma 2. (*Variational free energy*): If the autonomous dynamics of a particle or phenotype evince classical (Lagrangian) mechanics, then they can be expressed as minimising a variational free energy functional of Bayesian beliefs—about external states—encoded by their internal phenotypic states, $p_{\bar{\mu}}(\bar{\eta})$, under a generative model encoded by their (extended) genotype $p_{\bar{x}}(\bar{\eta}, \bar{\pi} | x_0)$:

$$\begin{aligned}\dot{\vec{\alpha}} &= \mathbf{D}\vec{\alpha} - \nabla_{\vec{\alpha}} F(\vec{\alpha}, \vec{s}) \\ F(\vec{\alpha}, \vec{s}) &= \underbrace{\mathbb{E}_{p_{\bar{\mu}}} [\mathcal{L}_{\bar{x}}(\bar{\eta}, \vec{s}, \vec{\alpha})]}_{\text{Energy constraint}} - \underbrace{\mathbb{E}_{p_{\bar{\mu}}} [\mathcal{L}_{\bar{\mu}}(\bar{\eta})]}_{\text{Entropy}} \\ &= \underbrace{D_{KL}[p_{\bar{\mu}}(\bar{\eta}) | p_{\bar{x}}(\bar{\eta} | x_0)]}_{\text{Complexity}} + \underbrace{\mathbb{E}_{p_{\bar{\mu}}} [\mathcal{L}_{\bar{x}}(\vec{s}, \vec{\alpha} | \bar{\eta})]}_{\text{— Accuracy}} \\ &= \underbrace{D_{KL}[p_{\bar{\mu}}(\bar{\eta}) || p_{\bar{x}}(\bar{\eta} | \vec{s}, \vec{\alpha}, x_0)]}_{\text{Divergence}} + \underbrace{\mathcal{L}_{\bar{x}}(\vec{s}, \vec{\alpha})}_{\text{— Log evidence}} \\ \mathcal{L}_{\bar{x}}(\vec{x}) &\triangleq -\ln p_{\bar{x}}(\vec{x} | x_0)\end{aligned}\tag{16}$$

This variational free energy can be rearranged in several ways. First, it can be expressed as an energy constraint minus the entropy of the variational density, which

licences the name free energy [93]. In this decomposition, minimising variational free energy corresponds to the maximum entropy principle, under the constraint that the expected Lagrangian is minimised [51,94]. The energy constraint is a functional of the marginal density over external and sensory states that plays the role of a generative model (i.e., parameterised by the extended genotype), namely, a joint density over causes (external dynamics) and their consequences (autonomous dynamics). Second—on a statistical reading—variational free energy can be decomposed into the (negative) log likelihood of particular paths (i.e., *accuracy*) and the KL divergence between posterior and prior densities over external paths (i.e., *complexity*). Finally, it can be written as the negative *log evidence* plus the KL divergence between the variational and conditional (i.e., posterior) density. In variational Bayesian inference [95], negative free energy is called an *evidence lower bound* or ELBO [96–98].

Proof. The sparse coupling—that underwrites a particular partition—means autonomous paths (i.e., generalised states) depend only on sensory paths. This means there is a (deterministic and injective) map from the most likely autonomous paths (of sufficiently high order generalised motion) to the conditional density over external paths, where both are conditioned on sensory paths. This injection means we can consider the conditional density over external paths as being parameterised by internal paths. We will call this a *variational density* (noting from (6) that internal paths are conditionally independent of external paths):

$$p_{\bar{\mu}}(\bar{\eta}) \triangleq p_{\bar{x}}(\bar{\eta} | \bar{s}, \bar{\mathbf{a}}, x_0) = p_{\bar{x}}(\bar{\eta} | \bar{s}, \bar{\mathbf{a}}, \bar{\mu}, x_0)$$

$$\begin{aligned} \bar{\alpha} &= \arg \min_{\bar{\alpha}} \mathcal{L}_{\bar{x}}(\bar{\alpha} | \bar{s}) \Rightarrow \\ \bar{\mu} &= \arg \min_{\bar{\mu}} \mathcal{L}_{\bar{x}}(\bar{\mu} | \bar{s}, \bar{\mathbf{a}}) \\ \bar{\mathbf{a}} &= \arg \min_{\bar{\mathbf{a}}} \mathcal{L}_{\bar{x}}(\bar{\mathbf{a}} | \bar{s}, \bar{\mu}) \end{aligned} \quad (17)$$

This definition means that the Lagrangian and variational free energy share the same minima, where their gradients vanish:

$$\begin{aligned} \bar{\alpha} &= \arg \min_{\bar{\alpha}} \mathcal{L}_{\bar{x}}(\bar{\alpha} | \bar{s}) = \arg \min_{\bar{\alpha}} F(\bar{\alpha}, \bar{s}) \\ \Rightarrow \nabla_{\bar{\alpha}} \mathcal{L}(\bar{\alpha}, \bar{s}) &= 0 \\ \Rightarrow \nabla_{\bar{\alpha}} F(\bar{\alpha}, \bar{s}) &= \nabla_{\bar{\alpha}} \underbrace{D_{KL}[p_{\bar{\mu}}(\bar{\eta}) || p_{\bar{x}}(\bar{\eta} | \bar{s}, \bar{\mathbf{a}})]}_{\text{Divergence}=0} + \underbrace{\nabla_{\bar{\alpha}} \mathcal{L}_{\bar{x}}(\bar{s}, \bar{\alpha})}_{=0} = 0 \end{aligned} \quad (18)$$

If autonomous dynamics are conservative, their trajectory is a path of least action and we can replace the Lagrangian gradients in (12) with variational free energy gradients to give (16). □

Remark 2. The free energy lemma (Lemma 2) associates negative fitness with variational free energy, such that phenotypic behaviour will appear to pursue paths of least free energy or greatest fitness. Because variational free energy is an upper bound on log evidence, the pursuit of maximum fitness can be read as self-evidencing [99]: namely, actively soliciting evidence for generative models endowed by evolution. In short, autonomous dynamics (appear to) actively infer external states under a generative model, whose parameters are (apparently) learned by minimising a path integral of variational free energy.

The functional form of variational free energy licences a teleological interpretation of autonomous dynamics; the internal paths can be read as the sufficient statistics or parameters of (approximate) Bayesian beliefs about external states, while active paths will

(appear to) change the posterior over external states to ‘fit’ internal (Bayesian) beliefs. In other words, active dynamics will look as if they are trying to fulfil the predictions of internal representations. A complementary interpretation inherits from the decomposition of variational free energy into complexity and accuracy. Minimising complexity means that generalised internal states encode Bayesian beliefs about external states that are as close as possible to prior beliefs, while generalised active states will look as if they are changing sensory states to realise those beliefs. These interpretations—in terms of *perception* and *action*—furnish an elementary but fairly expressive formulation of active inference. For example, the free energy formulations above have been used to emulate many kinds of sentient behaviour, ranging from morphogenesis [100], through action observation [101], to birdsong [102].

Although not developed here, the renormalisation group construction means that we can apply the same arguments to autonomous kinds—i.e., agents—at the slow scale. In other words, on average, the extended genotype of internal kinds comes to encode Bayesian beliefs about external kinds, while active kinds will look as if they are trying to realise those beliefs, via niche construction [77,103–105]. In virtue of the minimisation of variational free energy, we have an implicit maximum entropy principle, which brings us back to [21,22] via [49].

4. Variational Recipes

Effectively, we are describing the evolutionary developmental process with the following protocol:

- i. First, generate an ensemble of particles (i.e., extended phenotypes) by sampling their flow parameters and initial states (i.e., extended genotypes) from some initial density.
- ii. For each particle, find the path of least action using a generalised Bayesian filter (i.e., active inference).
- iii. After a suitable period of time, evaluate the path integral of variational free energy (i.e., action) to supply a fitness functional.
- iv. Update the flow parameters and initial states, using a stochastic gradient descent on the action (i.e., Darwinian evolution).

If this protocol were repeated for a sufficiently long period of time, it would converge to an attracting set, assuming this pullback attractor exists [32]. In statistical mechanics, this would be a nonequilibrium steady state, while in theoretical biology, it would correspond to an evolutionary steady state, at a certain timescale.

The notion of a steady state is clearly an idealization, as it assumes an unchanging environment. The local environments of all organisms are, however, moving targets, largely due to the activities of other organisms. Even if all of Life is considered a single population, it faces a changing local (i.e., biospheric) environment due to its—Life’s—own activities, as well as to bolide impacts and other abiotic causes. Hence, we can expect evolution to remain always ‘in process’ even for large, diverse populations. The assumption of an asymptotic evolutionary steady state is, therefore, effectively an assumption of a local (in time) steady state that has a lifetime long enough for evolutionary processes to be significant but short enough that the local environment of the evolving system can be considered approximately fixed. We now conclude with a simple application of the above protocol to a special case of selection in neurobiology.

A Numerical Study of Synaptic Selection

Figure 3 shows the results of a numerical study of selection processes, using the variational procedures above. This example illustrates the interplay between minimising variational free energy over somatic lifetimes and its use in selecting phenotypes at a slow, transgenerational, timescale. This example considers a relatively straightforward selection process in neurobiology, namely, synaptic selection in neurobiology, which illustrates the

nested scales over which free energy minimising processes evolve. Specifically, we simulated a single neuron (i.e., nerve cell) immersed in an environment constituted by potential pre-synaptic inputs in the surrounding neuropil. Unbeknown to the neuron (or more specifically, its dendritic tree), these presynaptic inputs fluctuated systematically with spatially structured waves of activation. These waves could only be detected by deploying postsynaptic specialisations (i.e., sensory states) in an ordered sequence along the dendrite. The details of this simulation are not important, and can be found in [106]. The key point here is that the cell's adaptive fitness—read as negative variational free energy—depends upon predicting its synaptic inputs through internal, intracellular dynamics that recapitulate the external, extracellular or environmental generation of sensory (synaptic) inputs. However, to do this, the dendrite has to have the right morphology, parameterised by the location of synapses on the dendritic surface.

To model learning and inference, the synapses were rendered more or less sensitive to their presynaptic inputs by optimising their sensitivity (a.k.a., precision) with respect to variational free energy in a biologically plausible fashion (i.e., using electrochemical equations of motion that performed a gradient flow on variational free energy). This meant that as the cell accumulated evidence from its presynaptic environment, its free energy decreased, and it became better at predicting its presynaptic inputs. However, this ability to predict depends upon selecting synapses that are located in the right order, along the dendrite.

To simulate synaptic selection, we used Bayesian model selection to compare the evidence for a cell's model with and without a particular synaptic connection. If the free energy increased, the postsynaptic specialisation was moved to another location at random. This process was repeated to simulate slow (Bayesian model) synaptic selection, until the phenotypic morphology of the dendrite was apt for accurately modelling (i.e., fitting) the waves of pre-synaptic input. In this example, the Bayesian model selection used Bayesian model reduction [107], based upon the optimised sensitivity (i.e., precision) of each synapse: very much along the lines of synaptic regression and implicit homeostasis [108–110]. Figure 3 shows the progressive reduction in free energy at a slow timescale as the synapses that enable the cell to predict or fit its environment are selected.




Figure 3. synaptic selection. This figure reports the results of numerical studies using fast free-energy minimising processes to model phenotypic dynamics and slow free-energy minimising processes to select phenotypic configurations or morphologies that, implicitly, have the greatest adaptive fitness or adapt to fit their environment. In this example, we focus on the selection of synapses of a brain cell (i.e., neuron) that samples presynaptic inputs from its neuropil (i.e., environment). The details of the generative model—used to simulate intracellular dynamics as a gradient flow on variational free energy—can be found in [107]. The key thing about these simulations is that—after a period of time—certain synapses were eliminated if Bayesian model selection suggested that their removal increased Bayesian model evidence (i.e., decreased variational free energy). (A): Findings in [111] suggest that neurons are sensitive to the pattern of synaptic input patterns. The image shows a pyramidal cell (blue) sampling potential presynaptic inputs from other cells (yellow) with postsynaptic specialisations (red). (B): In this model, pools of presynaptic neurons fire at specific times, thereby establishing a hidden sequence of inputs. The dendritic branch of the postsynaptic neuron comprises a series of segments, where each segment contains a number of synapses (here: five segments with four synapses each). Each of the 20 synapses connects to an axon of a specific presynaptic pool. These provide presynaptic (sensory) inputs at specific times over the length of a dendrite. If each of the 20 synapses were deployed in an orderly fashion across the five segments—as in the connectivity matrix—an orderly sequence of postsynaptic activations would be detected, and, implicitly. (C): The lower panels show the deployment of synaptic connections over 64 ‘generations’ (i.e., cycles), in which the precision (a.k.a. sensitivity) of synapses was used to eliminate synapses if they did not contribute to model evidence. Each ‘lifetime’ of the cell was 120 (arbitrary) time units, during which time two waves of activation were detectable. The upper panels score the ensuing increase in marginal likelihood or adaptive fitness (negative free energy) over the 64 generations. The left panel shows the accompanying increase in the sensitivity (i.e., log-precision) of the 20 synapses as they find the collective arrangement that maximises adaptive fit or model evidence for this (neuronal) environment.

5. Discussion

One insight from the above analysis is that populations are not necessarily quotient sets of equivalence classes. Put simply, there is no assumption that any given particle shares phenotypic or genotypic characteristics with any other particle. This observation is

interesting on two counts. First, it suggests that treating a population as an equivalence class of conspecifics may not be sufficient, in the sense that the population includes all of the (natural) kinds that interact to maintain their particular partition. The fact that all ‘individual’ multicellular eukaryotes appear to be holobionts—effectively, complex, multi-species ecosystems—bears this out [95,96]. The ‘genotype’ of such a system is a probability distribution of probability distributions, each of the latter over one of the component ‘species’ composing the holobiont. The phenotype of the holobiont, including its reproductive success and hence ‘fitness’ in the narrow reading, is a function of this bilevel probability distribution. Differential rates of genetic change between component genomes—and the fact that actions at the phenotypic level can alter the genotype as a probability distribution (e.g., humans can take anti- or probiotics)—complicate the difference in characteristic times assumed in Lemma 1, as discussed further below. Second, even if some agents share the same genotype, their phenotypes can specialise in distinct ways to minimise their joint variational free energies. This is obvious in the case of multicellular eukaryotes, all of which exhibit differentiation of cellular phenotypes during morphogenesis; see [89] for a worked example specifically employing the FEP formalism, and [97] for simulations demonstrating that multicellularity with differentiation provides a generic means of minimising VFE from the environment. These considerations together mandate a quintessentially co-evolutionary perspective that emphasises co-dependencies and co-creation [16,98–100].

However, the emergence of equivalence classes—e.g., ‘species’ of holobionts—begs explanation. A potential answer is the generalised synchrony between particles, as they find their joint variational free energy minima—and become mutually predictable; e.g., [52,91]. In an evolutionary setting, one can imagine the search for joint variational free energy minima leading to convergent evolution or speciation (Luc Ciompi, personal communication; [101]). Reproduction is, in all extant organisms, a matter of cell division, and closely related cells reap a free-energy advantage by working together [97]. An effective—though metabolically, morphologically, and behaviourally expensive—mechanism to protect this advantage is sex. The proliferation of species-specific morphological and behavioural specializations, together with the suppression of stem-cell pluripotency required to render sex obligate [102] in ‘higher’ eukaryotes, attests to the success of this strategy. From the present perspective, sex is a particularly elaborate feedback pathway—from the phenotypic to the genotypic scale—that preserves the integrity of the latter. It is, in other words, a mechanism that decreases VFE for the genome at the expense of increased VFE for the phenotype.

The synthesis of biological evolution and development on offer here is an example of a generalised synthesis: applicable, under the free energy principle, to all kinds of things. This synthesis can be read as generative models autopoietically generating entities and then using the ‘fit’ of the model to the niche as evidence for updating the model, in a cyclical process summarised in Figure 4.




Figure 4. Phylogeny and ontogeny as bottom-up and top-down causation.

6. Limitations

As with most applications of the free energy principle, the variational account alone does not supply a process theory. Rather, it starts from the assumption that a nonequilibrium (evolutionary) steady state exists and then describes the dynamics that the system must exhibit. Thus, the variational account enables various process theories to be proposed as specific hypotheses about biological systems. For example, the genetic variation in the above formulation follows from the Helmholtz decomposition or fundamental theorem of vector calculus. However, the ensuing stochastic gradient Langevin dynamics does not specify the particular processes that give rise to this kind of dynamics, e.g., [103]. There are many candidates one could consider: for example, simple rejection sampling or more involved genetic algorithms that provide a plausible account of bisexual reproduction [104,105]. A computationally expedient way of evaluating the requisite gradients—for example those for simulating artificial evolution—could call upon Bayesian model reduction [45,112]. Irrespective of the replication or reproduction process, it must, on the present analysis, conform to a stochastic gradient flow on ‘fitness’ with solenoidal mixing [72,78,79].

This openness to multiple process theories is an advantage of the current approach, both in convergence situations in which diverse genomes produce very similar phenotypes [112] and in the complementary situations in which a single genome supports diverse phenotypes. Neither situation is rare: genomes as different as those of *Amoeba proteus* and *Homo sapiens* can produce amoeboid cells, and the differentiated cells of any multicellular organism illustrate phenotypic diversity at the cellular level. While the general theory outlined here merely requires that some process exists, we can realistically expect one-to-many process mapping in both directions when dealing with real biological systems.

The primary offering of this variational formulation of natural selection—from an empirical perspective—is that one can hypothesise alternative forms for the Lagrangian. Each choice of Lagrangian will have consequences not only for the dynamics over physiological and developmental timescales but will also allow for predictions as to evolution over phylogenetic timescales. It is also worth noting that the account of natural selection set out here, in which genotypic evolution depends upon the action of phenotypic paths, applies to systems that satisfy the variational fitness lemma (Lemma 1): namely, the likelihood of an agent’s genotype corresponds to the likelihood of its phenotypic trajectory. While a plausible assumption—that is intuitively consistent with Darwinian evolution—we did not examine the conditions under which this assumption holds. This means there is an opportunity to further the ideas set out in this paper by examining the sorts of

stochastic systems in which the variational fitness lemma (Lemma 1) holds. It could be argued that Lemma 1 must hold at least in those systems where the genotype transforms into the phenotype retaining an equivalence within stochastic limits. For example, gene expression is the most fundamental level at which the genotype gives rise to the phenotype, and this mapping from genotype to phenotype is the subject of the many process theories studied by developmental biology. On a teleological view, one might further argue that active inference is necessary to maintain a high degree of equivalence during the course of this transformation and to preserve a correspondence between genotype and phenotype.

Available edge cases are, however, informative. Single mutations can induce saltatory changes in phenotype; a canonical example is the four-winged *Drosophila melanogaster* fly produced by combining three mutations, *abx*, *bx³*, and *pbx* of the bithorax complex in a single animal [113]. In complementary fashion, the planarian *Dugesia japonica* reproduces by fission followed by regeneration and has a heterogeneous, mixoploid genome with no known heritable mutants [114]; the phenotype of this animal has, however, remained stable for many thousands of generations in laboratories, and in all likelihood for millions of years in the wild. The phenotype can, moreover, be perturbed in saltatory fashion from one-headed to two-headed by an externally imposed bioelectric change; this altered phenotype is bioelectrically reversible but otherwise apparently permanent [115]. Engineering methods can create even more radically diverse phenotypes without genetic modifications, as demonstrated by the ‘xenobots’ prepared from *Xenopus laevis* skin cells, which adopt morphologies and behaviours completely unlike those that skin cells manifest when in the frog [116,117].

The availability of experimentally tractable edge cases of Lemma 1 provides an opportunity to further the ideas set out in this paper by examining the sorts of stochastic systems in which the variational fitness lemma (Lemma 1) holds. The kinds of edge cases mentioned above suggest, however, that Lemma 1 could be weakened to holding ‘up to’ saltatory events, including abiotic events such as bolide impacts, affecting genotype, phenotype, or both without substantially affecting the theory. Any systems that survive such events—any systems whose Markov blankets remain intact—simply carry on, undergoing learning, variation, and selection as if the saltatory event had never occurred.

One could suggest that Lemma 1, and the broader scope of the formalisms described here, may be applicable to systems where a population of entities engages in intergenerational replication (modelled here using the renormalisation operations), and where those entities at a faster timescale engage in rapid adaptation (e.g., development, learning, behaviour, modelled with active inference) during their lifetime. These two levels could, for example, model how genome-based intergenerational evolution sets initial conditions for organismal molecular and behavioural developments. For the faster intra-generational scale, the external states model the material basis of what the phenotype is a generative model of. For the slower inter-generational scale, the external states are updated through time as a process of renormalisation (reduction and grouping) of the extended genotype-phenotype.

7. Conclusions

This work attempts to unify the slow, multi-generational phylogenetic process of natural selection with the single-lifetime, phenotypic process of development (equations and notation summarized in Supplementary Materials). In this perspective, a bidirectional flow of information occurs as evolution imposes top-down constraints on phenotypic processes, and action selection provides evidence that is selected for by the environment (i.e., bottom-up causation). In this account, learning and inference occur through updating probabilistic beliefs via Bayesian model selection in evolutionary time and active inference in developmental time. The fitness of (extended) genotypes and (extended) phenotypes is selected for through the minimisation of the same free energy functional: Bayesian model evidence or marginal likelihood.

Further studies using both simulations and laboratory experiments are clearly needed to test this framework in the context of particular process theories that propose explicit functional connections between genotype and phenotype. While Lemma 1 is *prima facie* plausible in the case of idealised ‘central dogma’ organisms in which phenotype is largely determined by genotype within a tightly constrained, essentially static niche, the relation between genotype and phenotype in holobionts inhabiting realistic niches can be expected to be substantially more complex. ‘Egalitarian’ organisms, e.g., obligate symbionts or holobionts, comprising cells with different genotypes [118] and engineered systems—that offer cells radically different environments than they have experienced in phylogenetic evolution to date [119]—may be of particular interest for such studies.

Supplementary Materials: The following supporting information can be downloaded at: www.mdpi.com/10.3390/e25070964. A tabular glossary of mathematical expressions and technical terms used in this paper can be found at <https://www.activeinference.org/research/resources/variational-evolution>, hosted by the Active Inference Institute.

Author Contributions: All authors made substantial contributions to conception and design and writing of the article, and approved publication of the final version. All authors have read and agreed to the published version of the manuscript.

Funding: KF is supported by funding for the Wellcome Centre for Human Neuroimaging (Ref: 205103/Z/16/Z) and a Canada-UK Artificial Intelligence Initiative (Ref: ES/T01279X/1). DAF is supported by a National Science Foundation postdoctoral fellowship (ID 2010290).

Institutional Review Board Statement: No IRB approval was required or acquired for this study.

Informed Consent Statement: No ICS was required or acquired for this study.

Data Availability Statement: All data are included in the manuscript and Supplemental Materials.

Acknowledgments: We would like to thank members of the Active Inference Institute and their guests for invaluable discussions and guidance on the presentation of these ideas.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Darwin, C. *On the Origin of the Species by Natural Selection*; Murray: London, UK, 1859.
2. Dennett, D.C. *Darwin's Dangerous Idea: Evolution and the Meaning of Life*; Simon and Schuster: New York, NY, USA, 1996.
3. Fairbanks, D.J.; Abbott, S. Darwin's Influence on Mendel: Evidence from a New Translation of Mendel's Paper. *Genetics* **2016**, *204*, 401–405.
4. Mendel, G. Versuche über Pflanzen-Hybriden. 1866. Available online: <https://www.biodiversitylibrary.org/part/175272> (accessed on 14 June 2023).
5. Watson, J.D.; Crick, F.H.C. Molecular Structure of Nucleic Acids: A Structure for Deoxyribose Nucleic Acid. *Nature* **1953**, *171*, 737–738.
6. Dawkins, R.; Charles, S.; Dawkins, D.; Dawkins, R.A. *The Selfish Gene*; Oxford University Press: New York, NY, USA, 1989.
7. Noble, D. A theory of biological relativity: No privileged level of causation. *Interface Focus* **2012**, *2*, 55–64.
8. Keller, E.F. *The Mirage of a Space between Nature and Nurture*; Duke University Press: Durham, NC, USA, 2010.
9. Powell, J.R. Evolutionary Biology the Genetic Basis of Evolutionary Change. *BioScience* **1975**, *25*, 118. <https://doi.org/10.2307/1297112>.
10. Jablonka, E.; Lamb, M.J. *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*; Life and Mind: Philosophical Issues in Biology and Psychology Series; MIT Press: Cambridge, MA, USA, 2005. Available online: <https://psycnet.apa.org/fulltext/2005-04046-000.pdf> (accessed on 14 June 2023).
11. Heiner, M.; Gilbert, D. BioModel engineering for multiscale Systems Biology. *Prog. Biophys. Mol. Biol.* **2013**, *111*, 119–128.
12. Sultan, S.E. Eco-Evo-Devo. In *Evolutionary Developmental Biology: A Reference Guide*; Nuno de la Rosa, L., Müller, G.; Springer: Cham, Switzerland, 2017; pp. 1–13.
13. Smart, J.M. Evolutionary development: A universal perspective. In *Evolution, Development and Complexity*; Springer: Berlin/Heidelberg, Germany, 2019; pp. 23–92. https://doi.org/10.1007/978-3-030-00075-2_2.
14. Bruineberg, J.; Rietveld, E.; Parr, T.; van Maanen, L.; Friston, K.J. Free-energy minimization in joint agent-environment systems: A niche construction perspective. *J. Theor. Biol.* **2018**, *455*, 161–178.
15. Ellis, G.F.R.; Noble, D.; O'Connor, T. Top-down causation: An integrating theme within and across the sciences? *Interface Focus* **2012**, *2*, 1–3.

16. Carthey, A.J.R.; Gillings, M.R.; Blumstein, D.T. The Extended Genotype: Microbially Mediated Olfactory Communication. *Trends Ecol. Evol.* **2018**, *33*, 885–894.
17. Vanchurin, V.; Wolf, Y.I.; Katsnelson, M.I.; Koonin, E.V. Toward a theory of evolution as multilevel learning. *Proc. Natl. Acad. Sci. USA* **2022**, *119*, e2120042119. <https://doi.org/10.1073/pnas.2120037119>.
18. Vanchurin, V.; Wolf, Y.I.; Koonin, E.V.; Katsnelson, M.I. Thermodynamics of evolution and the origin of life. *Proc. Natl. Acad. Sci. USA* **2022**, *119*, e2120037119. <https://doi.org/10.1073/pnas.2120042119>.
19. McGee, R.S.; Kosterlitz, O.; Kaznatcheev, A.; Kerr, B.; Bergstrom, C.T. The cost of information acquisition by natural selection. *bioRxiv* **2022**. <https://doi.org/10.1101/2022.07.02.498577>.
20. Geisler, W.S.; Diehl, R.L. Bayesian natural selection and the evolution of perceptual systems. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **2002**, *357*, 419–448.
21. Campbell, J.O. Universal Darwinism As a Process of Bayesian Inference. *Front. Syst. Neurosci.* **2016**, *10*, 49.
22. Ramírez, J.C.; Marshall, J.A.R. Can natural selection encode Bayesian priors? *J. Theor. Biol.* **2017**, *426*, 57–66.
23. Attias, H. Planning by Probabilistic Inference. In Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics, Key West, FL, USA, 3–6 January 2003; pp. 9–16.
24. Botvinick, M.; Toussaint, M. Planning as inference. *Trends Cogn. Sci.* **2012**, *16*, 485–488.
25. Kaplan, R.; Friston, K.J. Planning and navigation as active inference. *Biol. Cybern.* **2018**, *112*, 323–343.
26. Millidge, B. Deep Active Inference as Variational Policy Gradients. *arXiv* **2019**, arXiv:1907.03876.
27. Price, G.R. Fisher’s “fundamental theorem” made clear. *Ann. Hum. Genet.* **1972**, *36*, 129–140.
28. Ao, P. Laws in Darwinian Evolutionary Theory. *arXiv* **2006**, arXiv:q-bio/0605020.
29. Frank, S.A. Natural selection. V. How to read the fundamental equations of evolutionary change in terms of information theory. *J. Evol. Biol.* **2012**, *25*, 2377–2396.
30. Parr, T.; Pezzulo, G.; Friston, K.J. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*; MIT Press: Cambridge, MA, USA, 2022.
31. Conant, R.C.; Ross Ashby, W. Every good regulator of a system must be a model of that system. *Int. J. Syst. Sci.* **1970**, *1*, 89–97.
32. Crauel, H.; Flandoli, F. Attractors for random dynamical systems. *Probab. Theory Relat. Fields* **1994**, *100*, 365–393.
33. Smith, J.M. Evolutionary game theory. *Physical D* **1986**, *22*, 43–49.
34. Kwon, C.; Ao, P. Nonequilibrium steady state of a stochastic system driven by a nonlinear drift force. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **2011**, *84*, 061106.
35. Schwabl, F. Phase Transitions, Scale Invariance, Renormalization Group Theory, and Percolation. In *Statistical Mechanics*; Schwabl, F., Ed.; Springer: Berlin/Heidelberg, Germany, 2002; pp. 327–404.
36. Cardy, J. *Scaling and Renormalization in Statistical Physics*; Cambridge University Press: Cambridge, UK, 1996.
37. Fields, C.; Glazebrook, J.F.; Levin, M. Minimal physicalism as a scale-free substrate for cognition and consciousness. *Neurosci. Conscious.* **2021**, *2021*, niab013.
38. Koide, T. Perturbative expansion of irreversible work in Fokker–Planck equation à la quantum mechanics. *J. Phys. A Math. Theor.* **2017**, *50*, 325001.
39. Haken, H. *Synergetics: An Introduction: Nonequilibrium Phase Transitions and Self-Organization in Physics, Chemistry and Biology*; Springer: Berlin/Heidelberg, Germany, 1978.
40. Buckminster Fuller, R. *Synergetics: Explorations in the Geometry of Thinking*; Estate of R. Buckminster Fuller: Sebastopol, CA, USA, 1982.
41. Hoeting, J.A.; Madigan, D.; Raftery, A.E.; Volinsky, C.T. Bayesian model averaging: A tutorial (with comments by M. Clyde, David Draper and E. I. George, and a rejoinder by the authors. *Stat. Sci.* **1999**, *14*, 382–417.
42. Gershman, S.J.; Niv, Y. Learning latent structure: Carving nature at its joints. *Curr. Opin. Neurobiol.* **2010**, *20*, 251–256.
43. Tenenbaum, J.B.; Kemp, C.; Griffiths, T.L.; Goodman, N.D. How to grow a mind: Statistics, structure, and abstraction. *Science* **2011**, *331*, 1279–1285.
44. Lu, H.; Rojas, R.R.; Beckers, T.; Yuille, A.L. A Bayesian Theory of Sequential Causal Learning and Abstract Transfer. *Cogn. Sci.* **2016**, *40*, 404–439.
45. Smith, R.; Schwartenbeck, P.; Parr, T.; Friston, K.J. An Active Inference Approach to Modeling Structure Learning: Concept Learning as an Example Case. *Front. Comput. Neurosci.* **2020**, *14*, 41.
46. Sella, G.; Hirsh, A.E. The application of statistical physics to evolutionary biology. *Proc. Natl. Acad. Sci. USA* **2005**, *102*, 9541–9546.
47. Jaynes, E.T. Information Theory and Statistical Mechanics. *Phys. Rev.* **1957**, *106*, 620–630.
48. Ramstead, M.J.D.; Sakthivadivel, D.A.R.; Heins, C.; Koudahl, M.; Millidge, B.; Da Costa, L.; Klein, B.; Friston, K.J. On Bayesian Mechanics: A Physics of and by Beliefs. *arXiv* **2022**, arXiv:2205.11543.
49. Friston, K. A free energy principle for a particular physics. *arXiv* **2019**, arXiv:1906.10184.
50. Kirchhoff, M.; Parr, T.; Palacios, E.; Friston, K.; Kiverstein, J. The Markov blankets of life: Autonomy, active inference and the free energy principle. *J. R. Soc. Interface* **2018**, *15*, 20170792.
51. Levin, M. The Computational Boundary of a “Self”: Developmental Bioelectricity Drives Multicellularity and Scale-Free Cognition. *Front. Psychol.* **2019**, *10*, 2688.
52. Palacios, E.R.; Razi, A.; Parr, T.; Kirchhoff, M.; Friston, K. On Markov blankets and hierarchical self-organisation. *J. Theor. Biol.* **2019**, 110089.

53. Parr, T.; Da Costa, L.; Friston, K. Markov blankets, information geometry and stochastic thermodynamics. *Philos. Trans. A Math. Phys. Eng. Sci.* **2020**, *378*, 20190159.
54. Fields, C.; Friston, K.; Glazebrook, J.F.; Levin, M. A free energy principle for generic quantum systems. *arXiv* **2021**, arXiv:2112.15242.
55. Graham, R. Covariant formulation of non-equilibrium statistical thermodynamics. *Z. Phys. B Condens. Matter* **1977**, *26*, 397–405.
56. Ao, P. Potential in stochastic differential equations: Novel construction. *J. Phys. A Math. Gen.* **2004**, *37*, L25.
57. Yuan, R.; Ma, Y.-A.; Yuan, B.; Ao, P. Potential Function in Dynamical Systems and the Relation with Lyapunov Function. *Chin. Phys. B* **2014**, *23*, 010505.
58. Friston, K.; Heins, C.; Ueltzhöffer, K.; Da Costa, L.; Parr, T. Stochastic Chaos and Markov Blankets. *Entropy* **2021**, *23*, 1220. <https://doi.org/10.3390/e23091220>.
59. Friston, K.; Da Costa, L.; Sajid, N.; Heins, C.; Ueltzhöffer, K.; Pavliotis, G.A.; Parr, T. The free energy principle made simpler but not too simple. *arXiv* **2022**, arXiv:2201.06387.
60. Nicolis, G.; Prigogine, I. *Self-Organization in Nonequilibrium Systems: From Dissipative Structures to Order Through Fluctuations*; Wiley: Hoboken, NJ, USA, 1977.
61. Seifert, U. Entropy production along a stochastic trajectory and an integral fluctuation theorem. *Phys. Rev. Lett.* **2005**, *95*, 040602.
62. Yan, H.; Zhao, L.; Hu, L.; Wang, X.; Wang, E.; Wang, J. Nonequilibrium landscape theory of neural networks. *Proc. Natl. Acad. Sci. USA* **2013**, *110*, E4185–E4194.
63. Jiao, S.; Xu, S.; Jiang, P.; Yuan, B.; Ao, P. Wright-Fisher dynamics on adaptive landscape. *IET Syst. Biol.* **2013**, *7*, 153–164.
64. Kerr, W.C.; Graham, A.J. Generalized phase space version of Langevin equations and associated Fokker-Planck equations. *Eur. Phys. J. B-Condens. Matter Complex Syst.* **2000**, *15*, 305–311.
65. Da Costa, L.; Friston, K.; Heins, C.; Pavliotis, G.A. Bayesian mechanics for stationary processes. *Proc. Math. Phys. Eng. Sci.* **2021**, *477*, 20210518.
66. Friston, K.; Stephan, K.; Li, B.; Daunizeau, J. Generalised Filtering. *Math. Probl. Eng.* **2010**, *2010*, 621670. <https://doi.org/10.1155/2010/621670>.
67. Laland, K.N.; Odling-Smee, F.J.; Feldman, M.W. Evolutionary consequences of niche construction and their implications for ecology. *Proc. Natl. Acad. Sci. USA* **1999**, *96*, 10242–10247.
68. Lehmann, L. The adaptive dynamics of niche constructing traits in spatially subdivided populations: Evolving posthumous extended phenotypes. *Evolution* **2008**, *62*, 549–566.
69. Welling, M.; Teh, Y.W. Bayesian Learning via Stochastic Gradient Langevin Dynamics. 2011. Available online: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.226.363> (accessed on 14 June 2023).
70. Ao, P. Emerging of Stochastic Dynamical Equalities and Steady State Thermodynamics from Darwinian Dynamics. *Commun. Theor. Phys.* **2008**, *49*, 1073–1090.
71. Seifert, U. Stochastic thermodynamics, fluctuation theorems, and molecular machines. *arXiv* **2012**, arXiv:1205.4176.
72. Zhang, F.; Xu, L.; Zhang, K.; Wang, E.; Wang, J. The potential and flux landscape theory of evolution. *J. Chem. Phys.* **2012**, *137*, 065102.
73. Lopes da Silva, F. Neural mechanisms underlying brain waves: From neural membranes to networks. *Electroencephalogr. Clin. Neurophysiol.* **1991**, *79*, 81–93.
74. Buzsáki, G.; Draguhn, A. Neuronal oscillations in cortical networks. *Science* **2004**, *304*, 1926–1929.
75. Lisman, J. Excitation, inhibition, local oscillations, or large-scale loops: What causes the symptoms of schizophrenia? *Curr. Opin. Neurobiol.* **2012**, *22*, 537–544.
76. Levin, M. Endogenous bioelectrical networks store non-genetic patterning information during development and regeneration. *J. Physiol.* **2014**, *592*, 2295–2305.
77. Manicka, S.; Levin, M. Modeling somatic computation with non-neural bioelectric networks. *Sci. Rep.* **2019**, *9*, 18612.
78. Ignacio-Espinoza, J.C.; Ahlgren, N.A.; Fuhrman, J.A. Long-term stability and Red Queen-like strain dynamics in marine viruses. *Nat. Microbiol.* **2020**, *5*, 265–271.
79. Schenk, H.; Schulenburg, H.; Traulsen, A. How long do Red Queen dynamics survive under genetic drift? A comparative analysis of evolutionary and eco-evolutionary models. *BMC Evol. Biol.* **2020**, *20*, 8.
80. Baross, J.A.; Martin, W.F. The Ribofilm as a Concept for Life's Origins. *Cell* **2015**, *162*, 13–15.
81. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444.
82. Feynman, R.P. Statistical mechanics: A set of lectures. notes taken by R. Kikuchi and H.A. Feiveson, edited by Jacob Shaham Feynman, P. Available online: <https://catalogue.nla.gov.au/Record/1947281/Details> (accessed on).
83. Sakthivadivel, D.A.R. A Constraint Geometry for Inference and Integration. *arXiv* **2022**, arXiv:2203.08119.
84. Beal, M.J. Variational Algorithms for Approximate Bayesian Inference. Ph.D. Thesis, University College London, London, UK, 2003. Available online: <https://discovery.ucl.ac.uk/id/eprint/10101435/> (accessed on 14 June 2023).
85. Winn, J.; Bishop, C.M. Variational Message Passing. *J. Mach. Learn. Res.* **2005**, *6*, 661–694.
86. Bishop, C.M. *Pattern Recognition and Machine Learning*; Springer: Berlin/Heidelberg, Germany, 2006.
87. Dauwels, J. On Variational Message Passing on Factor Graphs. In Proceedings of the 2007 IEEE International Symposium on Information Theory, Nice, France, 24–29 June 2007; pp. 2546–2550.
88. Hohwy, J. The self-evidencing brain. *Nous* **2016**, *50*, 259–285.

89. Friston, K.; Levin, M.; Sengupta, B.; Pezzulo, G. Knowing one's place: A free-energy approach to pattern regulation. *J. R. Soc. Interface* **2015**, *12*, 20141383.
90. Friston, K.; Mattout, J.; Kilner, J. Action understanding and active inference. *Biol. Cybern.* **2011**, *104*, 137–160.
91. Isomura, T.; Parr, T.; Friston, K. Bayesian Filtering with Multiple Internal Models: Toward a Theory of Social Intelligence. *Neural Comput.* **2019**, *31*, 2390–2431.
92. Laland, K.; Matthews, B.; Feldman, M.W. An introduction to niche construction theory. *Evol. Ecol.* **2016**, *30*, 191–202.
93. Constant, A.; Ramstead, M.J.D.; Veissière, S.P.L.; Campbell, J.O.; Friston, K.J. A variational approach to niche construction. *J. R. Soc. Interface* **2018**, *15*, 20170685.
94. Blackiston, D.; Lederer, E.; Kriegman, S.; Garnier, S.; Bongard, J.; Levin, M. A cellular platform for the development of synthetic living machines. *Sci. Robot.* **2021**, *6*, eabf1571. <https://doi.org/10.1126/scirobotics.abf1571>.
95. Guerrero, R.; Margulis, L.; Berlanga, M. Symbiogenesis: The holobiont as a unit of evolution. *Int. Microbiol.* **2013**, *16*, 133–143.
96. Bordenstein, S.R.; Theis, K.R. Host Biology in Light of the Microbiome: Ten Principles of Holobionts and Hologenomes. *PLoS Biol.* **2015**, *13*, e1002226.
97. Fields, C.; Levin, M. Somatic multicellularity as a satisficing solution to the prediction-error minimization problem. *Commun. Integr. Biol.* **2019**, *12*, 119–132.
98. Kauffman, S.A.; Johnsen, S. Coevolution to the edge of chaos: Coupled fitness landscapes, poised states, and coevolutionary avalanches. *J. Theor. Biol.* **1991**, *149*, 467–505.
99. Rosenman, M.; Saunders, R. Self-regulatory hierarchical coevolution. *Artif. Intell. Eng. Des. Anal. Manuf.* **2003**, *17*, 273–285.
100. Traulsen, A.; Claussen, J.C.; Hauert, C. Coevolutionary dynamics in large, but finite populations. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **2006**, *74*, 011901.
101. Kriz, J.; Tschacher, W. *Synergetik als Ordner. Die Strukturelle Wirkung der Interdisziplinären Ideen Hermann Hakens*; Pabst Science Publishers: Lengerich, Germany, 2017; pp. 15–20.
102. Fields, C.; Levin, M. Why isn't sex optional? Stem-cell competition, loss of regenerative capacity, and cancer in metazoan evolution. *Commun. Integr. Biol.* **2020**, *13*, 170–183.
103. Richardson, K. Genes and knowledge: Response to Baverstock, K. the gene an appraisal. *Prog. Biophys. Mol. Biol.* **2021**, *167*, 12–17. <https://doi.org/10.1016/j.pbiomolbio.2021.04.005>.
104. Casella, G.; Robert, C.P.; Wells, M.T. *Generalized Accept-Reject Sampling Schemes*; Institute of Mathematical Statistics Lecture Notes-Monograph Series; Institute of Mathematical Statistics: Beachwood, OH, USA, 2004; pp. 342–347.
105. Mehrabian, A.R.; Lucas, C. A novel numerical optimization algorithm inspired from weed colonization. *Ecol. Inform.* **2006**, *1*, 355–366.
106. Kiebel, S.J.; Friston, K.J. Free energy and dendritic self-organization. *Front. Syst. Neurosci.* **2011**, *5*, 80.
107. Friston, K.; Parr, T.; Zeidman, P. Bayesian model reduction. *arXiv* **2018**, arXiv:1805.07092.
108. Ravassard, P.; Pachoud, B.; Comte, J.C.; Mejia-Perez, C.; Scoté-Blachon, C.; Gay, N.; Claustrat, B.; Touret, M.; Luppi, P.H.; Salin, P.A. Paradoxical (REM) sleep deprivation causes a large and rapidly reversible decrease in long-term potentiation, synaptic transmission, glutamate receptor protein levels, and ERK/MAPK activation in the dorsal hippocampus. *Sleep* **2009**, *32*, 227–240.
109. Tononi, G.; Cirelli, C. Sleep function and synaptic homeostasis. *Sleep Med. Rev.* **2006**, *10*, 49–62.
110. Toutounji, H.; Pipa, G. Spatiotemporal computations of an excitable and plastic brain: Neuronal plasticity leads to noise-robust and noise-constructive computations. *PLoS Comput. Biol.* **2014**, *10*, e1003512.
111. Branco, T.; Clark, B.A.; Hausser, M. Dendritic discrimination of temporal input sequences in cortical neurons. *Science* **2010**, *329*, 1671–1675.
112. Conway Morris, S. Evolution: Like any other science it is predictable. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **2010**, *365*, 133–145.
113. Bender, W.; Akam, M.; Karch, F.; Beachy, P.A.; Peifer, M.; Spierer, P.; Lewis, E.B.; Hogness, D.S. Molecular Genetics of the Bithorax Complex in *Drosophila melanogaster*. *Science* **1983**, *221*, 23–29.
114. Fields, C.; Levin, M. Are Planaria Individuals? What Regenerative Biology is Telling Us About the Nature of Multicellularity. *Evol. Biol.* **2018**, *45*, 237–247.
115. Durant, F.; Morokuma, J.; Fields, C.; Williams, K.; Adams, D.S.; Levin, M. Long-Term, Stochastic Editing of Regenerative Anatomy via Targeting Endogenous Bioelectric Gradients. *Biophys. J.* **2017**, *112*, 2231–2243.
116. Kriegman, S.; Blackiston, D.; Levin, M.; Bongard, J. A scalable pipeline for designing reconfigurable organisms. *Proc. Natl. Acad. Sci. USA* **2020**, *117*, 1853–1859.
117. Kriegman, S.; Blackiston, D.; Levin, M.; Bongard, J. Kinematic self-replication in reconfigurable organisms. *Proc. Natl. Acad. Sci. USA* **2021**, *118*, e2112672118.
118. Strassmann, J.E.; Queller, D.C. The social organism: Congresses, parties, and committees. *Evolution* **2010**, *64*, 605–616.
119. Clawson, W.P.; Levin, M. Endless forms most beautiful 2.0: Teleonomy and the bioengineering of chimaeric and synthetic organisms. *Biol. J. Linn. Soc. Lond.* **2022**, blac073. <https://doi.org/10.1093/biolinnean/blac073>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.