



# The evolution of brain architectures for predictive coding and active inference

Giovanni Pezzulo, Thomas Parr, Karl Friston

## ► To cite this version:

Giovanni Pezzulo, Thomas Parr, Karl Friston. The evolution of brain architectures for predictive coding and active inference. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 2022, 377 (1844), 10.1098/rstb.2020.0531 . hal-04106061

**HAL Id: hal-04106061**

**<https://amu.hal.science/hal-04106061>**

Submitted on 30 May 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# The evolution of brain architectures for predictive coding and Active Inference

*Giovanni Pezzulo<sup>1</sup>, Thomas Parr<sup>2</sup>, Karl Friston<sup>2</sup>*

*1. National Research Council, Rome, Italy*

*2. Wellcome Trust, University College London, UK*

## Corresponding Author

Giovanni Pezzulo  
Institute of Cognitive Sciences and Technologies,  
National Research Council  
Rome, Italy  
[giovanni.pezzulo@istc.cnr.it](mailto:giovanni.pezzulo@istc.cnr.it)

## Abstract

This article considers the evolution of brain architectures for predictive processing. We argue that brain mechanisms for predictive perception and action are not late evolutionary additions of advanced creatures like us. Rather, they emerged gradually from simpler predictive loops (e.g., autonomic and motor reflexes) that were a legacy from our earlier evolutionary ancestors – and were key to solve their fundamental problems of adaptive regulation. We characterize simpler-to-more-complex brains formally, in terms of *generative models* that include predictive loops of increasing hierarchical breadth and depth. These may start from a simple homeostatic motif and be elaborated during evolution in four main ways: these include the *multimodal expansion* of predictive control into an allostatic loop; its *duplication* to form multiple sensorimotor loops that expand an animal's behavioural repertoire; and the gradual endowment of generative models with *hierarchical depth* (to deal with aspects of the world that unfold at different spatial scales) and *temporal depth*, (to select plans in a future-oriented manner). In turn, these elaborations underwrite the solution to biological regulation problems faced by increasingly more sophisticated animals. Our proposal aligns neuroscientific theorising—about predictive processing—with evolutionary and comparative data on brain architectures in different animal species.

**Keywords:** predictive processing; Active Inference; brain evolution; brain architecture; model selection; natural selection.

## Introduction

There is growing consensus that the brains of humans and other phylogenetically derived or advanced organisms operate in a predictive manner across perception (*predictive processing* [1,2]) and action control (*Active Inference* [3,4]). Yet, the ways in which our advanced predictive abilities may have arisen during evolution remains unclear. The goal of this article is to sketch an evolutionary history of brain architectures for predictive processing.

A central tenet of our proposal is that although prediction is often characterized as a complex cognitive function, it is not a late evolutionary addition of advanced animals like us. Rather, our complex predictive abilities (e.g., planning and imagination) emerged gradually (e.g. via phyletic gradualism or punctuated equilibrium) from simpler predictive and error-correction loops (e.g., motor and autonomic reflexes) that were already part of the brains of our earlier evolutionary ancestors – and were key to solving adaptive regulation problems [5–7].

We first consider the design of ancestral brain in terms of “generative models” that use simple predictive motifs for adaptive regulation. Then we discuss how structural designs were selected during evolution by differentiating generative models and by endowing them with hierarchical and temporal depth. In turn, these three expansions augmented the repertoire of species-specific behaviours and afforded increasingly more sophisticated predictive abilities, such as the planning and imagination, which characterise advanced animals.

### Predictive regulation and control are basic design principles of the brain

Every living organism faces a fundamental problem of adaptive regulation and control of its vital parameters, such as body temperature or glucose levels. Despite some living organisms (e.g., plants) survive without a brain, the possibility to solve adaptive control problems by moving around in the world – for example, to find food and shelter – might have been a strong selective pressure for the evolution of brains.

From a formal perspective, it is possible to conceptualize adaptive control as the imperative to “visit” only a limited set of possible states – or to remain in a (organism-dependent) niche that meets existential imperatives. One example of this is the imperative to keep body temperature around 37 degrees. In informational terms, the states that lie outside these acceptable boundaries are “surprising”. For example, a sensed body temperature much higher or lower than 37 degrees, or the sensation of being out of water for a goldfish, are “surprising” in this technical sense (technically, this kind of surprises is called surprisal or self-information). This implies that the imperative of living organism is to minimize the surprise of their sensory exchanges with the niche.

In turn, surprise is a function of two things: sensory *observations* sampled actively from the world and internally generated *predictions* about the observation. Specifically, surprise increases with the *discrepancy* between predictions and observations (under Gaussian assumptions, this discrepancy is the squared difference between predictions and observations), see Figure 1.

**Figure 1:** Action-perception cycle and predictive regulation. (A) Both perception and action reduce the discrepancy between predictions and observations, but the former operates by changing the organism’s beliefs (i.e., probability distributions) to make predictions closer to the observations and the latter operates by changing the niche, to make observations closer to predictions. (B) Active Inference formulates this process of discrepancy reduction as the minimization of variational free energy  $F$ , which is a function of approximate posterior beliefs  $Q$  and sensory observations  $y$ ; see [3]. The variational free energy  $F$  can be decomposed into two

*terms: divergence and evidence. The former is a Kullback-Leibler divergence between the approximate posterior  $Q(x)$  and the true posterior  $P(x|y)$ ; and the latter is the negative logarithm of the probability of observations or marginal likelihood  $P(y)$ . The figure highlights that changing beliefs corresponds to minimizing the first term (divergence) and changing observations corresponds to maximizing the second term (evidence).*

An organism can minimize this discrepancy in two ways: by *changing her mind* to predict better the next observation (perception) or by *changing the world* to ensure that the next observation complies with predictions (action), see Figure 1A. Active Inference formalises the ensuing resolution of surprise as the minimization of *variational free energy*: a quantity that is an upper bound on surprise and whose two components (divergence and surprise) map to perception and action, respectively; see Figure 1B and [3] for technical details.

As Figure 1 illustrates, action and perception operate synergistically to minimize the discrepancy between predictions and observations (or formally, *variational free energy*). Without perception, the animal would have trouble selecting appropriate actions (or have access only to a very limited repertoire, e.g., move faster if there is no food). Without acting to secure food, an animal would not survive for long to gather any more information or, more poetically, self-evidence [8].

Importantly, action is motivated by an “optimistic” belief about being able to realise one’s predictions. If a thirsty animal under the hot sun simply “changes her mind” and starts predicting sensations of sunburn and dehydration, she would not survive for long. Rather, the animal could predict a constant temperature of 37 degrees – even under the hot sun – and do her best to find shade. In other words, while the animal has to sometimes change its mind, it is endowed with some “priors” about the outcomes it prefers to solicit – e.g., the homeostatic prior of keeping body temperature at 37 degrees – that are of special importance and need to be protected from (excessive) revision. These priors play the same role as “set points” in cybernetics, to enable error-correction and (negative) feedback control [9–11].

One example of priors over preferred outcomes are “drive states”, such as hunger and thirst [12]. In Active Inference, these and other similar imperatives lie at the top of the animal’s control hierarchy and permit animals to continuously evaluate their current state in relation to their desired states of affairs; and to steer action if deviations are sensed [13]. To this aim, prior predictions from the top-down (e.g., that body temperature is expected to be at 37 degrees) are continuously compared with sensations from the bottom up (e.g., the interoceptive sensation of excessive body temperature) and any discrepancy calls for adaptive regulatory processes that resolve the discrepancy, thus implementing (negative) feedback control [14].

There are three important corollaries of this view. First, in Active Inference (and cybernetics) it is not an external stimulus but an internal event – a mismatch between a “prior prediction” (or “set point”) and a sensation – that triggers corrective actions. In other words, these theories describe animals as purposive entities driven by internally set goals, matching mechanisms and error-correction loops, not (only) by responses to environmental stimuli, as in theories that describe brains as a stimulus-response system. This is closely related to the good regulator theorem and perceptual control theory [10,15,16].

Second, the same mismatch (e.g., hyperosmolality) can engage multiple regulatory and error-correction processes, which range from simpler autonomic responses within the body (e.g., release of antidiuretic hormone) to complicated goal-directed plans that cancel errors through much longer loops through the environment (e.g., rent a beach umbrella or buy an ice cream on a sunny beach), depending on the allostatic complexity of the organism.

Third, the same error-correction scheme – that functions for internal (homeostatic and drive) states – can be used to control bodily movements and external variables, too. This is illustrated by the use of “equilibrium points” in motor control – and the fact that it is possible to guide bodily actions by setting fixed point attractors at desired movement endpoints [17]. Another illustration of the same principle is the notion of “goals” in cognitive theory [18]. In sophisticated animals like us, control can span cognitive and social goals, such as winning a tennis match or becoming member of a club. An animal will strive to achieve these cognitive goals through a similar error correction scheme as homeostatic control: by selecting courses of actions expected to reduce the discrepancy between the current and the goal state [19,20].

Given these three premises, we argue that simple error correction circuits, or basic “predictive motifs”, appeared very early in the evolutionary history of brain structure and functioned in largely similar ways across multiple domains (e.g., interoceptive control, motor control and goal-directed behaviour). The evolution of brain architecture proceeds by replicating the same scheme across multiple domains and behaviours (e.g., approach and escape; swimming, locomotion and climbing), thus forming multiple parallel sensorimotor loops. Furthermore, it proceeds by gradually augmenting the basic predictive motifs with more complex prediction and error correction circuits, which permit animals to extend their control from body physiology to the external (and social) environment. This analysis suggests that evolutionary history is constrained by the ecological pressures that animals need to face – and by the predictive motifs that are “evolvable” given the evolutionary trajectory to date.

### Formalizing brain design as structure learning in generative models

Before we sketch an evolutionary history of brain designs, we provide a formal way to interpret evolutionary trajectories, using the notion of *generative models*. Generative models are constructs from statistical theory that generate predictions about observations and are widely used in data science and machine learning. Here, we are concerned with generative models as explanations of how the brain works. Active Inference posits that creatures entail generative models of the hidden causes ( $x$ ) of their sensations ( $y$ ): see Figure 2 for a summary of this formalism. The brain’s models generate predictions (e.g., about environmental contingencies or the effects of body movements) and steers adaptive actions ( $u$ ) that bring about desired sensations (e.g., reach a food source).

**Figure 2.** *Generative model and generative process. This graph illustrates the notion of a generative model (in the brain) and a generative process (in the world) and their differences. Nodes correspond to probability distributions and edges to their statistical dependencies. Mathematically, a generative model may be formulated as the joint probability  $P(y,x)$  of observations  $y$  and the hidden states  $x$  of the world that generate these observations. The latter are referred to as hidden or latent states as they cannot be observed directly. This joint probability can be decomposed into two parts. The first is a prior  $P(x)$ , which denotes the organism's knowledge about the hidden states of the world, prior to seeing sensory data. The second is the likelihood  $P(y|x)$ , which denotes the organism's knowledge of how observations are generated from states. This graph also shows that the generative model and generative process are distinct. Both represent ways in which sensory data ( $y$ ) could be generated given hidden states ( $x$  or  $x^*$ ) and are represented through arrows from  $x$  or  $x^*$  to  $y$  to indicate causality. The difference is that the generative process is the ‘true’ causal structure by which data are generated, while the model is a construct used to draw inferences about the causes of data (i.e., use observations to derive ‘inferred states’). Hence,  $x$  and  $x^*$  can differ. Action ( $u$ ) is generated based upon the inferences made under a generative model. Action is shown here as part of the generative process, making changes to the world, despite being selected from the inferences drawn under the model.*

Generative models in the brain need not to be intended as internal copies of the external world. Their main goal is to control action, not (or not necessarily) to represent the external reality faithfully [21,22]. The hidden causes postulated in a generative model ( $x$ ) need not to be the same as the hidden states ( $x^*$ ) of the “true” process in the world (called generative process) that actually produce the organism’s sensations ( $y$ ); see Figure 2 for an illustration of this difference. In other words, the models we use to explain our sensorium may include hidden states that do not exist in the outside world, and vice versa. To the extent that an animal’s generative model is adequate to steer adaptive behaviour, the way it “carves nature” is not important.

As noted above, effective generative models also diverge systematically from environmental statistics. Most notably, some priors  $p(x)$  – such as the fact that body temperature fluctuates around 37 degrees – can be considered *prescriptive goals* that characterise the organism in terms of its prior preferences (“how the world should be”), rather than statements about the environment. While in standard Bayesian treatments priors are constantly updated during inference (following the motto that “today’s posterior is tomorrow’s prior”), in active inference some priors remain constant, because they have very high precision (and are therefore very recalcitrant to updating). In this way, they can function as set points that promote homeostasis. Nevertheless, these innate priors can be updated at slower (e.g., evolutionary) timescales [23], hence allowing organisms to “learn their preferences” and adapt to their niches.

Furthermore, generative models come in different varieties. This implies that the brains of different animals may correspond to different (simpler or more complex) generative models, which in turn enable different (more or less sophisticated) cognitive abilities. However, the diversity of generative models is not unbounded but has to follow two kinds of constraints. Firstly, all the generative models include some predictive motifs and – under gradualism – the most sophisticated generative models inherit (and extend) the predictive motifs of less sophisticated models. This implies that the space of the generative models that a particular animal can evolve is constrained by the generative models of its ancestors. Secondly, it is possible to define a priori which problems can or cannot be addressed by using different kinds of generative models. This implies that the generative model of each particular animal is tightly constrained by the statistics of their ecological niches and the control demands of their bodies [24,25].

Hence, it is possible to “reverse-engineer” the evolutionary history of brain designs, by (i) starting from simple “predictive motifs” putatively present in the brains of our evolutionary ancestors; then (ii) considering which expansions of these motifs (e.g., richer predictive loops) are possible and what novel biological problems they solve; and (iii) matching these basic or expanded motifs to the brains of specific animals in a lineage, by considering anatomy and the nature of the problems **and the natural statistics** of the animals’ niches.

In what follows, we provide examples of how the generative modelling perspective helps to scaffold an evolutionary trajectory of simpler-to-more-complex brain designs, in terms of generative models that include predictive loops of increasing complexity.

### Three examples of simple predictive motifs in ancestral brains

We start this tour through generative models using three examples of prediction and error correction loops. Given their simplicity, these generative models could be central to the design of “ancestral” brains.

#### *Generative models for the homeostatic control of interoceptive variables*

The generative models shown in Figure 3 afford the homeostatic (Figure 3A, “homeostat”) and allostatic (Figure 3C, “allostat”) regulation of a single interoceptive variable, which we call here “body temperature” for illustrative purposes; see [22] for a fully specified example. The (observable) variable  $y$  denotes the thermoreceptor activation and the (hidden) variable  $x$  denotes body temperature. The model infers the posterior belief over  $x$  by combining the prior belief over  $x$  and the interoceptive sensation  $y$ . The prior belief plays the role of a cybernetic set point to ensure that body temperature remains fixed at 37 degrees.

Much like a thermostat, this model maintains the requisite temperature by reporting the discrepancy between predicted and sensed thermoreceptor activation, given (Bayesian) beliefs about temperature, triggering an autonomic reflex ( $u$ ) resulting in, for example, vasodilatation, that resolves the prediction error. The functioning of this model can be appreciated by considering the (fictive) example of homeostatic control shown in Figure 3B. The figure plots the expectations of the prior and posterior, observations and autonomic actions over time. In this example,  $y$  is initially within an acceptable range but shortly afterwards it increases suddenly (say, following exposure to sunlight), causing the expected  $x$  under posterior beliefs to increase. The discrepancy between the predicted  $y$  given posterior beliefs and the measured  $y$  is registered as a prediction error. Given that the objective of active inference is to minimize prediction error the model triggers an autonomic action ( $u$ ), cancelling the prediction error. Note that in this example, the posterior belief over  $x$  is not a veridical representation of body temperature (it is “scaled down” by the prior belief). This exemplifies the fact that control demands are more important than representational accuracy.

**Figure 3.** Generative model for the allostatic regulation of a single interoceptive variable (in this example, body temperature). (A-B) Homeostat. This generative model includes an interoceptive thermoreceptor ( $y$ ) and a belief about body temperature ( $x$ ). Crucially, the prior over  $x$  is kept fixed and hence it acts as a cybernetic set point. Any discrepancy between the predicted thermoreceptor activity given beliefs about  $x$  and the measured  $y$  is registered as a prediction error that is cancelled out by an autonomic response ( $u$ ); for example, a thermoregulatory response. This is shown in an illustrative plot of the expectations of prior and posterior, observations and autonomic actions over time. (C) Allostat. This generative model extends the homeostat by including a second set of (exteroceptive) variables that correspond to light intensity ( $y_2$ ) and a belief about sunrise ( $x_2$ ). Furthermore, the model includes a predictive relationship between sunrise ( $x_2$ ) and body temperature ( $y$ ). In this way, inferring a sunrise can trigger the autonomic response ( $u$ ) of thermoregulation in an anticipatory manner, that is, before the sunlight actually increases body temperature. The upper parts of A and C are Bayesian networks, highlighting that  $y$  is conditionally dependent upon  $x$  with a directed arrow between the nodes (with more than one  $x$  and  $y$  in the model for the allostat). The lower parts show the form of neuronal message passing that could be used to solve these generative models. The red circles represent the expected values of  $x$ , which are used to make predictions about  $y$ . These are subtracted (red arrow with rounded end) from the measured  $y$  to form a prediction error (dark blue circle) which is used to update the expectation, and to drive action (light blue circle) that changes  $y$  such that the prediction error is resolved. Note the lateral modulatory connections in the allostatic network. See [26] for details.

This simulation illustrates a simple generative model (“homeostat”) supporting homeostasis via error correction: by registering prediction errors and actively cancelling them. This error correction scheme can be considered a basic “predictive motif” of generative models for Active Inference and is in continuity with feedback loops often found in the physiological control of body organs.

While for simplicity we considered a single homeostat that controls an aggregated variable (“body temperature”), regulatory problems—such thermoregulation—often imply the combination of



multiple, heterogeneous (e.g., feedback and feedforward) and partially independent mechanisms [27]. From this perspective, multiple homeostats that encode set points for separate (and simple) controlled variables may give rise to partially independent control loops. Furthermore, homeostats may operate synergistically with other (e.g., feedforward) mechanisms that implicitly promote the convergence of physiological variables to homeostatically valid values, rather than explicitly representing set points. Finally, more sophisticated mechanisms might also be at work. While in a homeostat the error is cancelled only when an expected (temperature) signal is sensed, physiological control can also act predictively (via the hypothalamus); for example, an animal can feel satiated after ingesting food, even before sugar becomes available. This example speaks to more prospective (or allostatic [28]) forms of regulation and control, which we discuss next.

### *Generative model for the allostatic control of interoceptive variables*

The “homeostat” is simple but limited: it can counter sensed changes of body temperature, but cannot anticipate predictable changes of body temperature (or other variables). In nature, there are several regularities (e.g., night-day or seasonal alternation) that can be easily incorporated to extend the above generative model as, technically speaking, empirical priors. The obvious advantage of predicting how our bodily and interoceptive variables will change is being able to exert some anticipatory (allostatic [28]) control.

For example, imagine that for an animal living in a hot zone, the pale light of the sunrise is predictive of the fact that its body temperature is going to become excessively high. If the agent's generative model incorporates this predictive relationship between sunrise and body temperature, it can anticipate changes in its (autonomic) state and initiate autonomic actions (e.g., vasodilatation) before the sun rises, pre-emptively mitigating the anticipated increase in temperature.

The generative model illustrated in Figure 3C (“allostat”) includes this predictive relation. Like the generative model of Figure 3A, the generative model of 3C includes thermoreceptors ( $y$ ) and the corresponding temperature ( $x$ ). However, it also includes a novel set of variables: photoreceptors ( $y_2$ ) and the light intensity from the sun ( $x_2$ ). The “allostat” is multimodal in the sense that it connects two sensory modalities: exteroceptive streams (e.g., photoreceptors) and interoceptive streams (e.g., thermoreceptors). The multimodal “allostat” (Figure 3C) could be realized gradually during evolution by slightly modifying the design of the unimodal “homeostat” (Figure 3A): namely, by including horizontal (predictive) relations between different modalities.

Critically, the two sets of variables of the “allostat” are coupled in the sense that “sunrise” is expected to cause both photoreceptor and thermoreceptor activation. By inscribing this causal structure into its generative model (i.e., neuronal networks), the “allostat” regulates body temperature in an anticipatory manner. As the animal photoreceptor activity increases, its expectation about the light intensity increases. In addition, and crucially, the animal anticipates a decrease in arteriolar tone *before* a thermoreceptor prediction error arises. By doing so, it prevents body temperature from increasing when the sun rises, avoiding the need for a homeostatic correction. This is in contrast to the functioning of the “homeostat”, where an autonomic action is only triggered *after* the thermoreceptor activity increases (see [29][22] for alternative formulations of allostatic control that appeal to hierarchical models; and [30] for a discussion of multimodal variables and convergence zones).

While we exemplified the “allostat” in an interoceptive regulation task, it can be applied more widely. For example, it can be used to model the predictive relations between distal senses (e.g., olfaction and vision) and proximity sensory (e.g., touch) and to allow animals to trigger escape behaviours when they see the shade of a predator. Below we provide a biological illustration of allostatic control of movement.

### ***Generative models for simple behavioural control***

The “homeostat” and the “allostat” permit the control of simple forms of swimming, locomotion, reaching, and other movements [31]. One biological example is provided by a zebrafish virtual reality study [32], which identified the neuronal underpinnings of error-correction during escape behaviour in the animal telencephalon: an evolutionarily homologous of brain circuits implied in action selection in other vertebrates (including mammals), such as the cortico-basal ganglia circuit [33].

In this study, animals were placed in a white start zone, which then become blue (aversive) or red (safe). If the start zone turned blue, the animal had to move forward to reach a red zone (GO trials). If the start zone turned red, the animal had to stay (NOGO trials). Failing to reach (or remain in) the safe zone implied an electric shock.

The study identified a neuronal ensemble that codes for a colour-based rule (i.e., “red is good”) and a separate neural ensemble (in one third of fish) that codes for a prediction error: a discrepancy between the predicted visual sensations of backward movement of the landscape during swimming and the actual visual input. Interestingly, the fish with the latter ensemble were more effective in avoidance behaviour. These fish may solve the task using a generative model analogous to the allostat of Figure 3C. The putative allostat could continuously generate motion predictions and prediction errors; and in turn, the fish could cancel the prediction errors by triggering appropriate (swimming or stopping) anticipatory responses, rather than (only) using autonomic reflexes. A crucial observation – to assess that behaviour was actually guided by prediction error minimization – was the fact that during control (open-loop virtual reality) trials in which swimming did not produce the predicted (backward movement) perceptions, the fish kept beating the tail – plausibly, because the prediction errors never disappeared.

This example suggests that simple regulatory circuits, such as escape circuits of the zebrafish, which are often considered quintessential examples of stimulus-response, could be supported – and improved – by error correction mechanisms [32]. The zebrafish telencephalon may generate various predictions (e.g., about expected backward movement of landscape) and use the corresponding prediction errors to guide active avoidance behaviour. The “allostat” architecture that we originally described in the context of interoceptive regulation is sufficient to solve the avoidance task of [32] (but it is possible that the zebrafish uses more sophisticated generative models that have temporal or hierarchical depth, see below).

### **An evolutionary algebra of structure learning**

Our central argument is that evolution proceeded via gradual elaborations of the “predictive motifs” illustrated above, under genetic constraints [34] and the selective pressure of novel problems to be solved, such as the control of more sophisticated bodies and the presence of richer ecological niches; e.g., when vertebrates begun to establish life on land some 400 million years ago.

Over successive generations, generative models can remain stable or be elaborated along four key dimensions, strongly limiting the space of “what is evolvable”. We have introduced the first kind of elaboration, from the (unimodal) homeostat to the (multimodal) allostat. A second kind of elaboration is the *duplication* of predictive motifs, which enlarges the animal's behavioural repertoire. The third and fourth dimensions equip the generative model with *temporal* and/or *hierarchical depth*, respectively. These two expansions enable richer predictive motifs that endow a cognitive sophistication, such as the possibility to plan or consider events that change on multiple timescales [35].

Figure 4 illustrates these four dimensions as if they were operations of an “evolutionary algebra” that defines the possible landscape of generative models (note that the algebra also includes an “identity” operation, to account for the fact that generative models can stay the same). In what follows, we discuss the duplicative and hierarchical operations – and their biological relevance.

**Figure 4.** *The five main dimensions of elaboration of generative models introduced in this paper, illustrated as operations of an “evolutionary algebra”. The identity (I) operation leaves the generative model as is. The duplication (I+I) operation replicates existing predictive motifs to form parallel sensorimotor loops. The allostast (A) operation endows the generative model with horizontal predictive relations between different modalities. The temporal depth (T) operation extends the generative models with separate variables for past, present and future states, hence affording prospective inference about the future (e.g., planning) and retrospective inference about the past. The hierarchical depth (H) operation extends the generative model with separate variables for states of affairs that change at different timescales (at faster timescales at the bottom levels and slower timescales at the higher levels), hence modelling narratives such as music and language, where nested timescales are relevant.*

### ***Duplicating predictive motifs enables multiple behaviours***

Generative models can expand by *duplicating* simple predictive motifs, to form a larger repertoire of species-specific behaviours, such as approach, avoidance, the control of the vibrissae [36] and visually guided grasping. The operator (I+I) in Figure 4 illustrates a generative model in which the same predictive motifs are duplicated and specialized, to form a “behaviour-based” architecture composed of multiple, parallel sensorimotor loops. Here, *duplication* means that the overall generative model comprises multiple smaller generative models (one for each ethologically valid behaviour) that operate in parallel and use the same error correction scheme. Instead, *specialization* means that the smaller generative models use different internal variables and are sensitive to different kinds of behaviour-related affordances and sensory information (e.g., space reachable with the arms for grasping or more distal space for locomotion).

Duplication may have been realized over evolution via the differentiation and parcellation of sensorimotor circuits that progressively acquire different (behaviour-specific) functional roles. For example, a single circuit for visually guided behaviour can be differentiated into two circuits, each specialized for a different kind of visually guided behaviour, such as escape versus foraging (see [37] for a detailed discussion of the differentiation between a retino-tectal circuit for spatial orientation and retino-telencephalic circuit for foraging). At the cortical level, this may imply the formation of different (behaviour specific) cortical fields and action maps, which may have more or less prevalence, depending on the relative importance (e.g., visual circuits may be more prevalent for diurnal than nocturnal animals) [38].

It is possible that early brain designs included generative models with multiple, replicated sensorimotor circuits. In keeping, Cisek [37] argued that “*the architecture of the early vertebrate could consist in a set of tectal circuits for different types of behaviors, each implemented as a closed feedback-control loop with the world.*” This organization is still visible in the organization of advanced brains, in which sensorimotor loops of various complexity are orchestrated by parallel basal ganglia-thalamocortical circuits [39] (as well as cerebellar loops). Further, in advanced brains, dorsomedial neocortex is organized around what Graziano calls *ethological action maps*: fronto-parietal circuits dedicated to different classes of species-specific actions [40].

Notably, an architecture composed of multiple sensorimotor loops requires additional mechanisms for their selection and prioritization. In the brain, behavioural selection (or more broadly, decision-making) is plausibly solved in a distributed manner, with different brain areas responsible for the most general to the most detailed aspects of the decision [37]. The hypothalamus may regulate general brain state (e.g., wakefulness or sleep), physiological cycles and basic allostasis. The arbitration between different classes of behaviours, such as approach or avoidance, may imply loops through the basal ganglia [41]. The arbitration between specific behaviours (e.g., different approach or grasping locations) could be solved as an “affordance competition” within the cortical system, at least in advanced animals [42]. Finally, ventromedial brain areas may support the prioritization of different hierarchical levels, depending on their current motivational value (see below our discussion of hierarchical models) [43,44]. Despite their diversity, these selection and prioritization circuits may all use similar dynamical competitive principles (e.g., biased competition [45]) – hence forming another “motif” of multipurpose architectures like brains.

From a structure learning perspective, duplication is an efficient way of building generative models: in the sense that the dynamics are conserved over different sensorimotor domains. This conservation is, mathematically, akin to factorising probability distributions in the generative model that has been discussed in terms of modular architectures [46] – and functional segregation as a principle of functional brain architectures [47–49]. In Bayesian statistics and physics, this kind of factorisation is ubiquitous and is known as a mean field approximation [50]. Indeed, the free energy bound on model evidence is defined in terms of a mean field approximation that affords an accurate and minimally complex explanation for (sensory) data [51].

### ***Endowing generative models with temporal depth supports prospective and retrospective inference***

The generative models discussed so far only consider present states and observations. However, they can be expanded into *temporally deep* models, whose variables explicitly represent future (and past) states and observations.

The operator  $T$  in Figure 4 illustrates a (discrete time) generative model having temporal depth. Temporally deep models support prospective and retrospective inference. For example, they permit predicting the short- or long-term consequences of actions (or action sequences, i.e., policies) and hence selecting the course of actions expected to deliver preferred outcomes, rather than just exploiting existing affordances [14]. Furthermore, they permit planning ahead, imagining novel situations and counterfactual reasoning, or the retrospective re-evaluation of one’s beliefs about the past in the light of novel evidence.

Various researchers have speculated that a major driving force for the development of deep temporal models was foraging. The increased cognitive and spatial demand of foraging – compared to visual orientation and landmark recognition – may have favoured the development of a (retino-telencephalic) foraging circuit distinct from another (retino-tectal) circuit for spatial orientation based on visible landmarks [37]. Intriguingly, the same hippocampal circuits that support spatial navigation and foraging are also involved in prospection and imagination [52]. This has led Buzsáki and Moser to propose that prospective functions have leveraged cognitive and predictive maps in the hippocampal-entorhinal system and hence “*mechanisms of memory and planning have evolved from mechanisms of navigation in the physical world*” [53].

The evolution of temporally deep models from simpler models could have been realized during evolution via the progressive parcellation of an initially undifferentiated model (i.e., a model that does not distinguish present from past and future) into a model that features separate latent states for the past, present and future. A key drive for this factorisation or parcellation may have been the

observation and progressive *internalization* of the sensorimotor sequences that the animal creates and experiences while acting – in other words, the self-modelling of one's own sequential behavioural patterns [14,54]; see [55] for a computational example. In turn, once these sensorimotor sequences are internalized to form a deep temporal model, they can be endogenously regenerated, to support memory and prospection. The progressive development of shallow models into deep temporal models may have occurred several times, across different brain structures, such as the hippocampus, the frontal cortex, and others – thus rendering advanced brains able to generate future predictions across various domains of cognition.

### ***Endowing generative models with hierarchical depth affords multi-scale inference***

So far, we have described generative models that can deal with aspects of the world that unfold at a single time scale. However, they can be expanded into *hierarchically deep* models, whose variables at different hierarchical levels encode latent states that unfold at different timescales. One example is a song: the melody remains the same even though the notes we hear (or sing) change rapidly. Similarly, a movie or narrative remains the same for several minutes, scenes remain the same for several seconds, but visual stimuli can change over hundreds of milliseconds.

The operator H in Figure 4 illustrates a generative model acquiring hierarchical depth. Such models permit modelling of narratives, songs, movies and other events that change at different temporal scales, by encoding variables that change more slowly (e.g., melodies or movies) at higher hierarchical levels and variables that change more rapidly (e.g., notes or visual scenes) at lower hierarchical levels. Two neurobiological examples of hierarchical organization are visual areas in mammals and the areas that control vocal gestures in birdsong [56,57].

It is possible that the same replication of sensorimotor circuits – across several domains (as discussed above) – also produced their “hierarchization” during evolution. In other words, different portions of sensorimotor circuits may have become specialized to deal with different timescales of action control: from simpler motor primitives to complex behaviours and finally to meaningful sequences that determine the temporal order of behaviours [58]. This hierarchical organization of behaviours (and brain structure) appears to be present in simple animals like the drosophila [59] and perhaps also in invertebrates, under the evolutionary pressure of dynamic foraging [60,61].

In more advanced animals, the hierarchical control of action may have expanded into sophisticated forms of “cognitive control” and “executive functions”, which help prioritize distal goals while inhibiting immediate affordances. It is possible to speculate that a capacity for cognitive control (associated to the development of prefrontal cortex [62]) could have evolved from the *internalization* of the control policies that animals used to realize goals in the external (and social) environment. In turn, this helped turn control from outside-in: from the control of states of the external environment to the self-regulation of own behaviour – which is the hallmark of executive function [63,64].

### **A phylogenetic tree of the evolution of generative models**

In the above, we formalized brain designs in terms of generative models that include predictive loops of various complexity; and then discussed the five main ways generative model designs can be elaborated – or the five main operations of an “algebra” of evolutionary structure learning (Figure 4).

This means that one can describe the evolutionary trajectory of brain designs in terms of a limited number of mutational operations over generative models. Figure 5 illustrates an example (portion of a) “phylogenetic tree” in which the application of mutational operations gives rise to a variety of generative models, marked with different coloured circles. While not all possible generative model

expansions are realized during evolution, some of them may correspond to actual brain designs of different animals. By prescribing which generative models are evolvable, in which branch of the phylogenetic tree they may occur and which adaptive problem they support, this kind of analysis may help align neuroscientific treatments of predictive processing with evolutionary and comparative data on brain architectures across species.

Interestingly, the mutational operators are commutative: the same generative model design can be obtained by executing the same operations in a different order. For example, both the generative models marked with green circles in Figure 5 result from the application of I+I, A, I, and T operators (albeit in a different order) and are therefore equivalent—note that the same is true of substitution mutations at a molecular level. The commutative property of mutational operators potentially sheds light on convergent evolution, and the process by which unrelated organisms evolve similar traits independently (and via different evolutionary histories), when they need to adapt to similar ecological niches.

***Figure 5.** Phylogenetic tree of generative model designs and putative correspondences with animal brains. Each branch of this example phylogenetic tree is generated by applying one of the four mutational operators discussed in Figure 4 (I: identity; I+I: duplication; A: allostatic (and multimodal) expansion; T: temporal depth; H: hierarchical depth). Two generative models are highlighted: a simpler generative model (blue circle), which may correspond to the brain of early vertebrates; and a more complex generative model (red circle), which may correspond to the brain of (some) primates. Note that the figure shows a reduced phylogenetic tree, for illustrative purposes: it is possible to expand the phylogenetic tree by applying the same operators ad libitum and in different orders. For the same reason, despite the generative model marked with the red circle appears at the apex of the example phylogenetic tree, it should be intended as an example and not as the endpoint of evolution.*

For illustrative purposes, we have highlighted two generative models in the phylogenetic tree. The former generative model (blue circle) may correspond to the brain of early vertebrates, as discussed above in relation to the zebrafish; and it would result from duplication, identity and allostatic operators. This generative model includes multiple parallel allostatic motifs for different behaviours (e.g., approach and avoidance) and an arbitration mechanism (e.g., subpallium, considered the zebrafish analogous of the mammalian basal ganglia [65]), not shown in the figure for simplicity.

The latter generative model (red circle) may correspond to the brain of (some) primates, and it would result from duplication, allostatic, temporal and hierarchical depth operators. This generative model includes several parallel sensorimotor loops for different behaviours; each having temporal and hierarchical depth. In this schematic, there are two parallel sets of hierarchical circuits, a dorsolateral control hierarchy and a ventromedial motivational hierarchy. In the dorsal hierarchy, lower-level circuits implement simple control loops that can be triggered directly by environmental affordances (e.g., food affords approaching), whereas higher-level circuits can contextualize lower-level circuits on the basis of distal goals, either to cooperatively support them or to override low-level affordances (e.g., a contextual memory of the presence of predators elevates an escape goal). Competition between levels is solved by a ventromedial motivational hierarchy, which prioritizes higher-level goals or lower-level affordances encoded in the dorsolateral control hierarchy, depending on which is expected to be more effective; see [43,44] for details.

This advanced (red) generative model has more sophisticated ways to perform biological regulation, compared to the simpler (blue) generative model. For example, while the simpler generative model can only avoid predators by escaping quickly, the more complex generative model can devise prospective strategies, such as building predator-proof econiches. However, it is important to remark that even in

the most complex organisms, evolutionarily novel solutions do not (completely) replace evolutionary older solutions to the same biological problems. Rather, older solutions largely remain available and can be selected in the appropriate conditions—thus rendering action selection context-sensitive [5,66]. For example, fear circuits appear to be organized hierarchically and to include simpler (reactive) and more sophisticated (cognitive) layers, which can be selected depending on external conditions, such as distance from the predator as well as internal conditions, such as arousal state [67].

## Discussion

In this article, we suggest that brain structure or design could be formalized as generative models; that the brain generative models of our evolutionary ancestors included simple “predictive motifs”; and that evolution proceeded via successive elaborations of these predictive motifs, into more complex architectures that we observe in advanced animals. We highlighted the four principal dimensions along which it is possible to advance simple predictive motifs into complex predictive motifs – and used this formal analysis to propose “phylogenetic trees” of generative models and corresponding brain designs. While the evolutionary trajectory of designs for predictive processing proposed here is certainly tentative and incomplete, we consider it a first step towards the alignment of predictive brains and evolutionary studies of neuroanatomy in different species.

The view that error correction mechanisms can encompass both simpler (homeostatic and allostatic) and more sophisticated (cognitive) forms of adaptive behaviour differs significantly from prevalent perspectives in psychology and neuroscience, which tend to appeal to separate sets of mechanisms for sensorimotor processing and simple cognition (reactive controllers) and for higher cognition (predictive controllers), respectively. A widespread perspective—inherited from behaviourism—is that the building blocks of adaptive behaviour in the brain are stimulus-response mappings (or policies) slowly reinforced by rewards, which in their most minimal form, might correspond to links between sensory and motor neurons that arose early during evolution [68]. The brain can leverage cognitive sophistication and depth by allowing the above stimulus-response rules to be selectively activated by higher (control) pathways—that develop much later during evolution and which may serve predictive regulation [69]. Another influential (“dual theory”) perspective assumes that the brain uses two entirely separate mechanisms for decision-making, one intuitive and one deliberative [70] (which are sometimes equated with distinct, model-free and model-based controllers of reinforcement learning [71]). In these and other, related views (e.g., the “triune brain” idea [72]), mechanisms for higher cognition (including predictive loops) are late evolutionary inventions. However, it is questionable whether the segregations implied in “dual theory” and related views can be reconciled with the gradualism and continuity of evolutionary processes [34,37] – and how mechanisms for higher cognition might have emerged during evolution on top of mechanisms for sensorimotor processing, if the two are supposed to be completely different from one another.

The alternative hypothesis pursued here is that early brain designs already used error correction mechanisms, as envisaged by cybernetic theories such as the *Test-Operate-Test-Exit* (TOTE) model [73] and Active Inference. One benefit of this hypothesis is that it does not require any evolutionary jumps to explain cognitively sophisticated behaviour. **Notably, closed loop control in the brain may not be a recent evolution invention: it may have originated** from even simpler forms of metabolic control and the synthesis of nutrients within a cell membrane; and then expanded outside it to realize what we call “behaviour” [16,74,75] – hence realizing a continuity of life and cognition.

We have considered how generative models could have expanded but we did not address “why”. Although not our focus here, the process by which selective pressure generates new neuronal structures can itself be expressed as a surprise or free energy minimising process [76,77]. This follows because Fisher’s fundamental theorem can be read as Bayesian belief updating, based upon marginal likelihood or model evidence [78]. Specifically, one can cast natural selection as nature’s way of

performing Bayesian model selection – sometimes known as structure learning [79] (i.e., learning the right kind of generative model for this econiche) – where model evidence just is adaptive fitness.

This perspective (which is still speculative and not unchallenged [80]) suggests that the complexity of the ecological niche determines the level of complexity the brain needs to have, in order to be (Bayes) optimal. Statistically, maximising model evidence is equivalent to maximising the difference between accuracy and complexity. Hence, any increase in model complexity over evolution is (only) licensed by a greater increase in model accuracy (where model evidence corresponds to accuracy minus complexity). In other words, brains only increase their complexity with sufficient ecological demands, e.g., the animal has to control a complex body or deal with complex situations with the requisite degree of accuracy. This is because having a more complicated brain does not help if you live in a simple niche. This is why evolution is not necessarily a linear road towards increasingly more complex brains, unless those brains make the niche more complicated. However, as the niche comes more complex – in virtue of being populated by increasingly complex specifics and their niche construction – there is a necessary increase in the complexity of the generative models apt for predicting those niches accurately. For example, the “social brain” hypothesis states that the necessity to predict and deal with the sophisticated social dynamics was a main driver of the evolution of large brains and sophisticated cognitive abilities in our species [81]. In short, the gradualism expressed as a progressive increase in complexity rests on the circular causality implicit in the modelling of an eco-niche that is itself constituted – and constructed – by increasingly complicated phenotypes [82–85].

The balancing of accuracy and complexity might be at work not just between species but also between different mechanisms within the same brain. Above, we discussed two examples – thermoregulation and fear circuits – in which multiple heterogeneous mechanisms (e.g., feedforward and feedback; simpler homeostats and more complex models having temporal depth), which plausibly reached maturity at different instants during evolution, coexist and solve control problems synergistically. These mechanisms might specialize to deal with distinct aspects of the organism’s econiche, following the principle that the recruitment of a more complex mechanism is (only) licensed by the necessity to deal with a more challenging problem, which would render simpler mechanisms inaccurate. This (bounded rational) perspective helps explaining in which sense it is optimal to match cognitive sophistication to task demands. Notably, this perspective also harmonizes simple (e.g., stimulus-response) controllers popular in optimal control and reinforcement learning within the “predictive processing” view offered in this paper. For example, a simulative study shows that both reactive and predictive controllers can be expressed within an overarching generative modelling perspective; and that active inference prioritizes the simpler reactive controller over the more complex predictive controllers, in specific circumstances that render the former sufficiently accurate (e.g., when contextual uncertainty has been reduced) [86].

One objection that could be raised — to the account on offer here—is that there is a simpler form of control in multicellular systems that is not (obviously) predictive. This is the 2-neuron reflex, comprising a sensory and a motor neuron [68]. When the quantity sensed by the former deviates from some target value (i.e., set point) the sensory neuron increases or decreases its firing rates relative to its resting activity. This prompts an increase or decrease in the activity of the motor neuron, until the sensed quantity is returned to its target value. Superficially, this example seems to be a purely feedforward mechanism (sensory to motor), with no predictive loop. However, an alternative interpretation of this is as a system whose model provides a context-free prediction that the sensed quantity is at its set point. The loop from sensed quantity to sensory neuron to motor neuron to sensed quantity could then be seen as a feedback loop. The implication is that the 2-neuron reflex is in fact the simplest form of predictive control. Why is this a useful perspective? First, it answers the objection outlined above. Secondly, and more interestingly, it furnishes a point of continuity with more complex models. Including an interneuron between the two, and allowing for this 3-neuron pathway to be modulated by other



neural systems, gives us a context sensitive prediction of the sort shown throughout the models illustrated in this paper. This continuity is essential from an evolutionary perspective, where we need to account for incremental steps in which one predictive mechanism may be built upon another.

Finally, it is important to acknowledge that the [co]evolution of brain designs, bodies, ecological and cultural niches are interdependent. Given that here we were interested in the evolution of brain designs, we assumed a brain-centric perspective and conveniently focused on generative models in the animal's brain. However, cognition does not need to be confined "in the skull" but can be extended outside it, to cover (for example) tools and social dimensions [87]. Furthermore, body design (and not just brain design) plays an important role in solving control problems [88]. **Acknowledging that cognition can be "extended" and "embodied" suggests that not all aspects of control need to be solved by (or represented in) a central generative model.** For example, active inference emphasizes that some aspects (that roughly correspond to the notion of inverse models in optimal control) can be directly implemented by reflex arcs in the spinal cord [89]. Other aspects of control might be offloaded to appropriately designed niches (e.g., turnarounds to coordinate multiple drivers) or exploit convenient characteristics of body design (e.g., the prehensility of hands) to alleviate the burden of brain generative models. These and other examples (see [90]) show that considering extended and embodied aspects of cognition provides a broader view of adaptive behaviour, in which a generative modelling capability may be distributed across the brain, body and environmental niche. In turn, applying the generative modelling perspective across brains, bodies and environments could help understand their synergistic interactions during evolution.

## Acknowledgments

This research received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement Nos. 785907 and 945539 (Human Brain Project SGA2 and SGA3) to GP, the European Research Council under the Grant Agreement No. 820213 (ThinkAhead) to GP. This work was conducted under funding for the Wellcome Centre for Human Neuroimaging (Ref: 205103/Z/16/Z)

## References

1. Rao RP, Ballard DH. 1999 Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci* **2**, 79–87. (doi:10.1038/4580)
2. Friston K. 2005 A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci* **360**, 815–836. (doi:10.1098/rstb.2005.1622)
3. Friston K, FitzGerald T, Rigoli F, Schwartenbeck P, Pezzulo G. 2017 Active Inference: A Process Theory. *Neural Comput* **29**, 1–49. (doi:10.1162/NECO\_a\_00912)
4. Parr T, Pezzulo G, Friston KJ. 2022 *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press.
5. Pezzulo G, Rigoli F, Friston KJ. 2015 Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*.
6. Barrett LF, Simmons WK. 2015 Interoceptive predictions in the brain. *Nat Rev Neurosci* **16**, 419–429. (doi:10.1038/nrn3950)
7. Seth AK, Friston KJ. 2016 Active interoceptive inference and the emotional brain. *Phil. Trans. R. Soc. B* **371**, 20160007.
8. Hohwy J. 2016 The Self-Evidencing Brain. *Noûs* **50**, 259–285. (doi:10.1111/nous.12062)
9. Ashby WR. 1952 *Design for a brain*. Oxford, England: Wiley.
10. Powers W. 1973 *Behavior, the control of perception*. Aldine de Gruyter.
11. Wiener N. 1948 *Cybernetics: or Control and Communication in the Animal and the Machine*. Cambridge, MA: The MIT Press.
12. Hull CL. 1943 *Principles of Behaviour*. New York: Appleton-Century-Crofts.
13. Chanes L, Barrett LF. 2016 Redefining the Role of Limbic Areas in Cortical Processing. *Trends Cogn. Sci. (Regul. Ed.)* **20**, 96–106. (doi:10.1016/j.tics.2015.11.005)
14. Pezzulo G, Cisek P. 2016 Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition. *Trends in Cognitive Sciences* (doi:10.1016/j.tics.2016.03.013)
15. Conant RC, Ashby WR. 1970 Every good regulator of a system must be a model of that system. *Intl. J. Systems Science* , 89–97.
16. Ashby WR. 1952 *Design for a brain*. Oxford, England: Wiley.
17. Feldman AG. 2009 New insights into action–perception coupling. *Exp Brain Res* **194**, 39–58. (doi:10.1007/s00221-008-1667-3)
18. Pezzulo G, Castelfranchi C. 2009 Thinking as the Control of Imagination: a Conceptual Framework for Goal-Directed Systems. *Psychological Research* **73**, 559–577.
19. Pezzulo G, Rigoli F, Friston KJ. 2015 Active Inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology* **136**, 17–35.
20. O'Reilly RC, Hazy TE, Mollick J, Mackie P, Herd S. 2014 Goal-Driven Cognition in the Brain: A Computational Framework. *arXiv:1404.7591 [q-bio]*
21. Tschantz A, Seth AK, Buckley CL. 2020 Learning action-oriented models through active inference. *PLOS Computational Biology* **16**, e1007805. (doi:10.1371/journal.pcbi.1007805)
22. Tschantz A, Barca L, Maisto D, Buckley CL, Seth AK, Pezzulo G. 2021 Simulating homeostatic, allostatic and goal-directed forms of interoceptive control using Active Inference. *bioRxiv* , 2021.02.16.431365. (doi:10.1101/2021.02.16.431365)
23. Sajid N, Tigas P, Zakharov A, Fountas Z, Friston K. 2021 Exploration and preference satisfaction trade-off in reward-free learning. *arXiv:2106.04316 [cs, q-bio]*
24. Berkes P, Orbán G, Lengyel M, Fiser J. 2011 Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* **331**, 83–87.

(doi:10.1126/science.1195870)

25. Pezzulo G, Zorzi M, Corbetta M. 2020 The secret life of predictive brains: what's spontaneous activity for? (doi:10.31234/osf.io/qus3h)
26. Friston KJ, Parr T, de Vries B. 2017 The graphical brain: Belief propagation and active inference. *Network Neuroscience* **1**, 381–414. (doi:10.1162/NETN\_a\_00018)
27. Romanovsky AA. 2007 Thermoregulation: some concepts have changed. Functional architecture of the thermoregulatory system. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology* **292**, R37–R46. (doi:10.1152/ajpregu.00668.2006)
28. Sterling P. 2012 Allostasis: a model of predictive regulation. *Physiology & behavior* **106**, 5–15.
29. Stephan KE *et al.* 2016 Allostatic Self-efficacy: A Metacognitive Theory of Dyshomeostasis-Induced Fatigue and Depression. *Front Hum Neurosci* **10**. (doi:10.3389/fnhum.2016.00550)
30. Meyer K, Damasio A. 2009 Convergence and divergence in a neural architecture for recognition and memory. *Trends in Neurosciences* **32**, 376–382. (doi:10.1016/j.tins.2009.04.002)
31. Friston K *et al.* 2012 Dopamine, Affordance and Active Inference. *PLoS Comput Biol* **8**, e1002327. (doi:10.1371/journal.pcbi.1002327)
32. Torigoe M, Islam T, Kakinuma H, Fung CCA, Isomura T, Shimazaki H, Aoki T, Fukai T, Okamoto H. 2021 Zebrafish capable of generating future state prediction error show improved active avoidance behavior in virtual reality. *Nat Commun* **12**, 5712. (doi:10.1038/s41467-021-26010-7)
33. Mueller T, Wullimann MF. 2009 An evolutionary interpretation of teleostean forebrain anatomy. *Brain Behav Evol* **74**, 30–42. (doi:10.1159/000229011)
34. Krubitzer L. 2007 The magnificent compromise: cortical field evolution in mammals. *Neuron* **56**, 201–208. (doi:10.1016/j.neuron.2007.10.002)
35. Friston K, Da Costa L, Hafner D, Hesp C, Parr T. 2021 Sophisticated Inference. *Neural Comput* **33**, 713–763. (doi:10.1162/neco\_a\_01351)
36. Ahissar E, Kleinfeld D. 2003 Closed-loop Neuronal Computations: Focus on Vibrissa Somatosensation in Rat. *Cereb Cortex* **13**, 53–62. (doi:10.1093/cercor/13.1.53)
37. Cisek P. 2019 Resynthesizing behavior through phylogenetic refinement. *Atten Percept Psychophys* (doi:10.3758/s13414-019-01760-1)
38. Krubitzer L. 2009 In search of a unifying theory of complex brain evolution. *Ann N Y Acad Sci* **1156**, 44–67. (doi:10.1111/j.1749-6632.2009.04421.x)
39. Alexander GE, DeLong MR, Strick PL. 1986 Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci* **9**, 357–381. (doi:10.1146/annurev.ne.09.030186.002041)
40. Graziano MS. 2016 Ethological Action Maps: A Paradigm Shift for the Motor Cortex. *Trends in cognitive sciences* **20**, 121–132.
41. Redgrave P, Prescott TJ, Gurney K. 1999 The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* **89**, 1009–1023.
42. Cisek P. 2007 Cortical mechanisms of action selection: the affordance competition hypothesis. *Phil. Trans. R. Soc. B.* **362**, 1585–1599.
43. Kouneiher F, Charron S, Koechlin E. 2009 Motivation and cognitive control in the human prefrontal cortex. *Nat Neurosci* **12**, 939–945. (doi:10.1038/nn.2321)
44. Pezzulo G, Rigoli F, Friston KJ. 2018 Hierarchical active inference: a theory of motivated control. *Trends in cognitive sciences* **22**, 294–306.
45. Desimone R, Duncan J. 1995 Neural mechanisms of selective visual attention. *Annu Rev Neurosci* **18**, 193–222. (doi:10.1146/annurev.ne.18.030195.001205)
46. Parr T, Sajid N, Friston KJ. 2020 Modules or mean-fields? *Entropy* **22**, 552.

47. Ungerleider LG, Haxby JV. 1994 “What” and “where” in the human brain. *Current Opinion in Neurobiology* **4**(2), 157–65.
48. Zeki S. 2005 The Ferrier Lecture 1995 behind the seen: the functional specialization of the brain in space and time. *Philosophical Transactions of the Royal Society B: Biological Sciences* **360**, 1145–1183.
49. Friston K, Buzsáki G. 2016 The functional anatomy of time: what and when in the brain. *Trends in cognitive sciences* **20**, 500–511.
50. Kschischang FR, Frey BJ, Loeliger H-A. 2001 Factor graphs and the sum-product algorithm. *IEEE Transactions on information theory* **47**, 498–519.
51. Winn J, Bishop CM, Jaakkola T. 2005 Variational message passing. *Journal of Machine Learning Research* **6**.
52. Szpunar KK, Spreng RN, Schacter DL. 2014 A taxonomy of prospection: introducing an organizational framework for future-oriented cognition. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 18414–18421. (doi:10.1073/pnas.1417144111)
53. Buzsáki G, Moser EI. 2013 Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat Neurosci* **16**, 130–138. (doi:10.1038/nn.3304)
54. Buzsáki G, Peyrache A, Kubie J. 2014 Emergence of Cognition from Action. *Cold Spring Harb. Symp. Quant. Biol.* **79**, 41–50. (doi:10.1101/sqb.2014.79.024679)
55. Stoianov I, Maisto D, Pezzulo G. 2020 The hippocampal formation as a hierarchical generative model supporting generative replay and continual learning. *bioRxiv* , 2020.01.16.908889. (doi:10.1101/2020.01.16.908889)
56. Yu AC, Margoliash D. 1996 Temporal Hierarchical Control of Singing in Birds. *Science* **273**, 1871–1875. (doi:10.1126/science.273.5283.1871)
57. Friston K, Kiebel S. 2009 Cortical circuits for perceptual inference. *Neural Netw* **22**, 1093–1104. (doi:10.1016/j.neunet.2009.07.023)
58. Tinbergen N. 1951 *The study of instinct*. New York: Oxford University Press.
59. Berman GJ, Bialek W, Shaevitz JW. 2016 Predictability and hierarchy in *Drosophila* behavior. *PNAS* **113**, 11943–11948. (doi:10.1073/pnas.1607601113)
60. Stout D. 2010 The evolution of cognitive control. *Top Cogn Sci* **2**, 614–630. (doi:10.1111/j.1756-8765.2009.01078.x)
61. Hills TT. 2011 The evolutionary origins of cognitive control. *Top Cogn Sci* **3**, 231–237. (doi:10.1111/j.1756-8765.2011.01135.x)
62. Passingham RE, Wise SP. 2012 *The Neurobiology of the Prefrontal Cortex: Anatomy, Evolution, and the Origin of Insight*. 1st edn. Oxford University Press, USA.
63. Barkley RA. 2001 The executive functions and self-regulation: An evolutionary neuropsychological perspective. *Neuropsychology Review* **11**, 1–29.
64. Pezzulo G, Castelfranchi C. 2009 Thinking as the Control of Imagination: a Conceptual Framework for Goal-Directed Systems. *Psychological Research* **73**, 559–577.
65. Cheng R-K, Jesuthasan SJ, Penney TB. 2014 Zebrafish forebrain and temporal conditioning. *Philosophical Transactions of the Royal Society B: Biological Sciences* **369**, 20120462. (doi:10.1098/rstb.2012.0462)
66. Verschure P, Pennartz CMA, Pezzulo G. 2014 The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Philos Trans R Soc Lond B Biol Sci.* , 369: 20130483.
67. Qi S, Hassabis D, Sun J, Guo F, Daw N, Mobbs D. 2018 How cognitive and reactive fear circuits optimize escape decisions in humans. *Proc Natl Acad Sci U S A* **115**, 3186–3191. (doi:10.1073/pnas.1712314115)
68. Horridge, G. 1968 The origins of the nervous system. In *The structure and function of nervous tissue*, NY:Acadeimc Press.
69. Miller EK, Cohen JD. 2001 An integrative theory of prefrontal cortex function. *Annu*

- Rev Neurosci* **24**, 167–202. (doi:http://dx.doi.org/10.1146/annurev.neuro.24.1.167)
70. Kahneman D. 2003 A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist* **58**, 697–720.
  71. Daw ND, Niv Y, Dayan P. 2005 Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience* **8**, 1704–1711. (doi:10.1038/nn1560)
  72. MacLean PD. 1990 *The triune brain in evolution: Role in paleocerebral functions*. Springer Science & Business Media.
  73. Miller GA, Galanter E, Pribram KH. 1960 *Plans and the Structure of Behavior*. New York: Holt, Rinehart and Winston.
  74. Maturana HR, Varela FJ. 1980 *Autopoiesis and Cognition: the Realization of Living*. Dordrecht, Holland: D. Reidel Pub.
  75. Cisek P. 1999 Beyond the computer metaphor: Behavior as interaction. *Journal of Consciousness Studies* **6**, 125–142.
  76. Hesp C, Ramstead M, Constant A, Badcock P, Kirchhoff M, Friston K. 2019 A multi-scale view of the emergent complexity of life: A free-energy proposal. In *Evolution, Development and Complexity*, pp. 195–227. Springer.
  77. Badcock PB, Friston KJ, Ramstead MJ. 2019 The hierarchically mechanistic mind: A free-energy formulation of the human psyche. *Physics of life Reviews* **31**, 104–121.
  78. Frank SA. 2012 Natural selection. V. How to read the fundamental equations of evolutionary change in terms of information theory. *Journal of evolutionary biology* **25**, 2377–2396.
  79. Campbell JO. 2016 Universal Darwinism as a process of Bayesian inference. *Frontiers in Systems Neuroscience* **10**, 49.
  80. Colombo M, Wright C. 2021 First principles in the life sciences: the free-energy principle, organicism, and mechanism. *Synthese* **198**, 3463–3488. (doi:10.1007/s11229-018-01932-w)
  81. Dunbar RIM. 1998 The social brain hypothesis. *Evolutionary Anthropology: Issues, News, and Reviews* **6**, 178–190. (doi:10.1002/(SICI)1520-6505(1998)6:5<178::AID-EVAN5>3.0.CO;2-8)
  82. Laland K, Matthews B, Feldman MW. 2016 An introduction to niche construction theory. *Evolutionary ecology* **30**, 191–202.
  83. Bruineberg J, Rietveld E, Parr T, van Maanen L, Friston KJ. 2018 Free-energy minimization in joint agent-environment systems: A niche construction perspective. *Journal of theoretical biology* **455**, 161–178.
  84. Constant A, Ramstead MJD, Veissière SPL, Campbell JO, Friston KJ. 2018 A variational approach to niche construction. *Journal of The Royal Society Interface* **15**, 20170685. (doi:10.1098/rsif.2017.0685)
  85. Veissière SP, Constant A, Ramstead MJ, Friston KJ, Kirmayer LJ. 2020 Thinking through other minds: A variational approach to cognition and culture. *Behavioral and Brain Sciences* **43**.
  86. Friston K, FitzGerald T, Rigoli F, Schwartenbeck P, O’Doherty J, Pezzulo G. 2016 Active inference and learning. *Neuroscience & Biobehavioral Reviews* **68**, 862–879. (doi:10.1016/j.neubiorev.2016.06.022)
  87. Nave K, Deane G, Miller M, Clark A. 2020 Wilding the predictive brain. *WIREs Cognitive Science* **11**, e1542. (doi:10.1002/wcs.1542)
  88. Pfeifer R, Bongard JC. 2006 *How the Body Shapes the Way We Think*. MIT Press.
  89. Adams RA, Shipp S, Friston KJ. 2013 Predictions not commands: active inference in the motor system. *Brain Struct Funct* **218**, 611–643. (doi:10.1007/s00429-012-0475-5)
  90. Clark A. 2015 *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*.

Oxford University Press, Incorporated.