Hedonic Foraging: From Reward to Action

Ana Clemente[1]* and **Olivier Penacchio**[2]*

[1] Department of Cognitive Neuropsychology, Max Planck Institute for Empirical Aesthetics (Germany), ana.clemente@ae.mpg.de, https://www.aesthetics.mpg.de/institut/mitarbeiterinnen/ana-clemente.html
[2] Bridging Research in AI and Neuroscience, Computer Vision Center, Computer Science Department, Universitat Autònoma de Barcelona (Spain), openacchio@gmail.com, https://openacchio.github.io/
* Equal contribution / Shared first authorship

Short abstract (≤ 100 words)

Assigning hedonic value shapes behaviour across organisms, yet how these values translate into motivated behaviour remains unclear. We introduce *hedonic foraging* as a mechanism by which hedonic evaluation guides behaviour via iterative reward maximisation. This framework integrates hedonic evaluation with active inference, whereby expected free energy encodes wanting—the motivational drive toward rewarding outcomes—while variational free energy corresponds to liking—the hedonic impact of the outcome. Hedonic foraging unifies the mechanisms underlying behaviours ranging from allostasis to art appreciation, casting cognition and motivated behaviour as purposive engagement with the environment and its affordances.

Long abstract (≤ 250 words)

Hedonic evaluation shapes behaviour across organisms. Yet, the mechanisms that transform hedonic values into motivated behaviour remain poorly understood, and key questions remain unanswered. How do habits, wanting, and liking or disliking interact to generate behaviour? Do common principles govern behaviour across species and cognitive systems, from basic survival mechanisms to complex cultural endeavours? We propose a unifying theory of motivated behaviour that integrates hedonic evaluation with active inference—a framework in which perception, action and cognition emerge from the imperative to minimise free energy. We argue that variational free energy provides a natural measure of liking and disliking, while expected free energy encodes wanting—the motivational drive toward rewarding outcomes. This synthesis, which we term *hedonic foraging*, offers a principled explanation for how organisms navigate their environments based on pleasure-driven learning and decision-making. It posits that motivated behaviour emerges from the fundamental pursuit of maximising reward, with hedonic evaluation motivating and monitoring actions through wanting and liking mechanisms. We demonstrate that hedonic foraging captures a wide range of behaviours, from allostatic processes common to all animals to sophisticated human activities such as art appreciation and cultural exploration. Crucially, our integrative perspective fosters mutual enrichment between hedonic evaluation and active inference—two fundamental approaches to behaviour. Ultimately, hedonic foraging reinforces the view that cognition is essentially about purposive engagement in adaptive interactions with the environment and its affordances and offers a unified framework for investigating cognition and behaviour.

Keywords (5–10)

Active inference, Bayesian inference, choice, hedonic evaluation, hedonic value, learning, liking, predictive coding, preference, wanting

1. Introduction

Motivated behaviour is usually examined through highly specialised lenses, focusing on idiosyncratic biological adaptations, neural processes and cognitive mechanisms. This leads to a proliferation of theories specific to particular fields or taxa, with limited prospects for unification. Yet, one of the main endeavours of cognitive psychology and neuroscience is to unveil general mechanisms of behaviour (Berridge, 2004; Friston, 2010), raising the fundamental question: Is there a common mechanism driving behaviour across the spectrum of complexity and species? For instance, could the same principles explain a *C. elegans* orienting in a nutrient gradient, an endotherm foraging to maintain metabolic balance and a human formulating complex plans to achieve distant goals?

When searching for unifying theories in biology and the behavioural sciences, evolution always emerges as a key principle (Gintis, 2007; Mesoudi et al., 2006). To be biologically meaningful, behaviour must be adaptive, enhancing the organism s chances of survival and reproduction by improving its interaction with the environment. Behavioural mechanisms are shaped by the need to ensure ontogenetic and phylogenetic success (Duckworth, 2009; Snell-Rood, 2013).

Adaptive behaviour relies on evaluating the environment—its elements, relationships and affordances (Chiel & Beer, 1997; Schwarz, 2007). The fundamental neurobiological process of assessing stimuli as beneficial or harmful, rewarding or punishing, is referred to as *hedonic evaluation.* This process underpins the capacity to compare alternatives, make decisions and prioritise actions (Berridge & Kringelbach, 2015; Rosenthal, 2017). Whether selecting food, choosing a mate or engaging in leisure activities, organisms with a reward system are driven by the pursuit of rewards and avoidance of threats (Berridge & Kringelbach, 2008; Friston, 2013; Kringelbach, 2005; Rangel et al., 2008). This reward-seeking behaviour has been fundamental to brain evolution (Schultz, 2015). Since the emergence of chordates (Cisek, 2022), neurons in the reward system compute value signals by integrating sensory input with interoceptive and executive information, reflecting the organism's past experiences, states, goals and expectations (Rangel et al., 2008), thereby motivating and evaluating behaviour (Berridge & Kringelbach, 2013, 2015).

Motivated behaviour unfolds in three interrelated phases, each involving specific mechanisms and value signals (Berridge & Kringelbach, 2013; Pessiglione & Lebreton, 2015). In the appetitive phase, *wanting* mechanisms anticipate the reward of behavioural outcomes, guiding approach or avoidance, biasing perception towards salient features and modulating neural activity in perceptual and cognitive systems. During the *consummatory* phase, these motivational aspects integrate with executive and motor systems to implement behavioural choices. Finally, in the satiety phase, *liking* (and *disliking*) mechanisms evaluate the behavioural outcomes, signalling reward and informing future behaviour through learning (Dayan, 2022).[1] Hedonic values are thus affective states that reward or punish behavioural decisions and inform future behaviour (Berridge & Kringelbach, 2015).

Yet, the mechanisms by which the information conveyed by hedonic values is transformed into behaviour are not fully understood. For example, how does the pleasurable taste of a fruit urge

[1] The processes and mechanisms discussed and proposed in this paper are not necessarily explicit or conscious, although they may be accompanied by explicit or conscious experience. Accordingly, we primarily refer to (implicit) core *wanting* (incentive salience) and *liking* (hedonic impact), as opposed to (explicit) conscious wanting (cognitive incentive) and liking (conscious pleasure) (Berridge & Dayan, 2021).

someone to buy it *and* someone else to cultivate an orchard? Whereas the former seeks immediate, simple reward, the latter engages in a long-term, risky and costly endeavour aiming for a future, complex reward. The picture gets even more complicated when multiple hedonic values conflict, as when animals explore new environments to seek novel and better resources despite facing a greater risk of predation. Hedonic evaluation explains the neuropsychological mechanisms of reward, but it does not explain how hedonic values are computed and translated into concrete actions. As a result, it fails to make specific predictions about future behaviour. This limitation arises because hedonic evaluation lacks a computational mechanism for selecting and implementing behavioural choices and, therefore, fails to establish causal relationships between hedonic value and actual behaviour.

The most primary behaviours are rooted in the biological imperative to maintain homeostasis, driven by evaluative processes that compare the organism's state with reference values (Cannon, 1929). This imperative provides a natural link to action, as deviations from homeostatic set points generate hedonic signals that motivate corrective behaviours to restore balance (Ramsay & Woods, 2014). For instance, organisms require precise glycaemia ranges, experiencing pleasure when approaching them (e.g., eating) and distress when deviating from them (e.g., feeling hungry or experiencing indigestion). Importantly, the behaviour to approach such preferences depends on the organism and the context of evaluation, as eating may be pleasurable when hungry but unpleasant when satiated (Berridge, 1991; Skov, 2023). Therefore, quantitative measures capturing the role of hedonic value in motivating behaviour, and thus a computational framework for motivated behaviour, must be grounded in comparisons with preferences and tailored to embodied, situated behaviour (Roth & Jornet, 2013; Shapiro, 2011).

The imperative of maintaining homeostasis has long been recognised as central for living organisms, with the view that deviations trigger corrective mechanisms through reflexive actions (Cannon, 1929). However, recent approaches focus on the anticipatory nature of physiological regulation, a process called *allostasis* (Kleckner et al., 2017; Sterling, 2012). Allostasis implies a crucial shift from a reactive to a *predictive* perspective on the brain, because a predictive brain anticipates future events, objects and demands (Clark, 2013). This has two key implications: First, the brain projects its preferred future outcomes, generating internal representations of the consequences of possible actions (Conant & Ross Ashby, 1970; Pezzulo & Cisek, 2016). Second, the brain commands anticipatory actions, namely those that, according to its internal model, lead to its preferred outcomes (Wiener, 1948).

These two implications call for an *enactive* approach whereby the brain does not merely store internal models of the world to infer the causes of incoming information but, more fundamentally, develops models that enable a dynamic interaction between the organism and its environment. Such an interaction entails estimating conditions and affordances upon which it may *act* to fulfil its needs and goals. This perspective brings together two pivotal heritages in conceptualising and understanding perception, and the very *raison d'être* of perception itself (Friston, 2010; Pezzulo et al., 2024). First, in Helmholtz's view (von Helmholtz, 1866), perception is a process of inference of the causes of the incoming sensory input. The brain constantly combines top-down information from an internal generative model and bottom-up sensory evidence in order to estimate what is going on in the environment. This view has been framed using a proper probabilistic setting in the *Bayesian brain hypothesis* (Doya et al., 2007; Gregory, 1980; Pouget et al., 2013; Rao & Ballard, 1999). Second, in Gibson's (Gibson, 1986) ecological theory of perception, the brain's primary function is to perceive affordances, i.e., chances for action in the environment. In this action-oriented perspective, the brain is directly

tuned to what the environment can offer to the organism without requiring complex internal processing.

To meaningfully explain and predict motivated behaviour, a mechanistic framework must therefore integrate predictive and enactive perspectives on the brain, proposing a probabilistic brain inference account of affordances in Bayes-optimal behaviour (Friston et al., 2006; Kording & Wolpert, 2006). Moreover, it must adopt an evolutionary perspective grounded in the simple axiom that an organism's behaviour serves the imperative to survive in its environment (Friston, 2010). This axiom implies that organisms maintain life by adopting adaptive behaviours that bring them to preferred states, reflecting optimal living conditions within their ecological niche (Friston et al., 2012; Pezzulo et al., 2015).

Active inference provides a mathematical framework to operationalise such a process (Friston, 2010). Specifically, in active inference, behaviour reflects the imperative to *minimise free energy*, a proxy for the divergence between actual and preferred states (Friston, 2010; Pezzulo et al., 2024). In this article, we demonstrate how a broad range of behaviours can be accounted for by integrating hedonic evaluation with the core constructs of active inference, thus providing a biologically plausible computational framework for investigating behaviour. We demonstrate the potential of the proposed unified framework by modelling behaviour along a continuum, from basic allostatic processes shared by all organisms to more complex activities typically deemed uniquely human, such as art appreciation. This proposal allows us to bridge the gap and clarify the causal relationship between fundamental physiological processes involved in hedonic evaluation and action selection, offering a cohesive explanation for the motivational underpinnings of behaviour across different levels of complexity. By framing hedonic evaluation within active inference, we can better understand how organisms, from simple life forms to humans, navigate their environments, make decisions and pursue goals in a way that balances immediate rewards with long-term survival and well-being.

We start by deriving a correspondence between the main components of hedonic evaluation, namely hedonic impact (core *liking* or *disliking*) and incentive salience (core *wanting*) (Berridge & Robinson, 1998), and the three components that guide policy choice in active inference, namely *free energy*, *expected free energy* and *habi*t (K. Friston et al., 2017). From this correspondence, we propose a mechanism by which hedonic evaluation drives behaviour: *hedonic foraging*. Following this exposition, we showcase hedonic foraging with a series of example scenarios to illustrate its suitability to account for a continuum of behaviours. We argue that hedonic foraging ascribes complex cognitive functions such as appreciating art to the set of behaviours already accounted for by active inference, therefore welding core biological mechanisms and sophisticated cultural pursuits. We next discuss hedonic foraging in relation to existing proposals and the implications and contributions of the framework. To conclude, we emphasise the consequences and potential of a unifying, enactive and embodied computational mechanistic framework for understanding cognition and behaviour.

2. Hedonic foraging: Hedonic evaluation as active inference

As previously outlined, the processes by which hedonic evaluation, a core motivational process, translates into a framework for understanding and predicting behaviour remain undefined. In particular, there is no clear mechanism for computing distinct hedonic values, integrating them and guiding behaviour over time. To fill this fundamental gap and understand motivated behaviour, we propose mapping hedonic evaluation onto the constructs of active inference, a principled theory that explains and predicts embodied cognition and behaviour (Friston, 2010; Pezzulo et al., 2024). The central move is to express the role of hedonic evaluation in monitoring and motivating free energy minimisation.

Free energy minimisation is at the core of active inference. The theory stems from the truism that organisms must maintain their bodily parameters within some limits to stay alive. This pursuit is framed using the concept of *surprise*[2] (Friston, 2010; Pezzulo et al., 2024). Surprising observations (e.g., exteroceptive, interoceptive, proprioceptive) are those that deviate from reference observations, which are called *preferences* and reflect the range of values that are characteristic and beneficial for an organism's functioning (Friston, 2010). In active inference, preferences are modelled as *expectations* about observations reflecting states of the world aligned with the organism's needs and goals. Surprise thus corresponds to *prediction errors* when such expectations are not met. Active inference posits that a single, unified principle drives behaviour: the imperative of minimising surprise, or, more precisely, minimising a tractable proxy of surprise called *free energy* (Friston, 2010; Pezzulo et al., 2024) (**Box 1**).

An organism minimises free energy through a process of inference based on a *generative model*, namely a probabilistic mapping between hidden states of the world s and their associated observations o. The (true) states of the world are called *hidden* because the organism may not have direct access to the true unfolding process—the *generative process*—occurring in the world (**Box 1**). Instead, the organism relies on observations to infer the underlying states. Following the tradition of perception as inference—originating with Helmholtz and further developed in the Bayesian brain hypothesis (Doya et al., 2007; Gregory, 1980; Rao & Ballard, 1999; von Helmholtz, 1866)—, the inference process is conceived as an inversion of the generative model, estimating the most likely states of the environment given observations.

Active inference extends Bayesian theories by incorporating *action* into the inference process. In this framework, the brain does not merely infer causes passively; it also generates actions to minimise discrepancies between predicted and actual observations. To minimise surprise and maintain their preferred states, organisms continuously invert their generative model, comparing predicted observations with actual observations, to actively reduce the gap between expected and observed outcomes. This surprise-minimisation process drives both perception—updating beliefs as in the tradition of perception as inference (von Helmholtz, 1866)—and action—modifying the environment to achieve preferred outcomes (Friston, 2010; Gallagher, 2017). Consequently, active inference is inherently enactive, incorporating action into the inference process to enable adaptive behaviour in a dynamic world (Friston, 2010; Pezzulo et al., 2024; Ramstead et al., 2020). For example, organisms cannot survive outside a specific range of glycaemic levels, and deviation with respect to their preferred (physiologically ideal) range

[2] The concept of surprise in active inference comes from information theory and does not necessarily align with the concept of psychological surprise.

constitutes a surprising state. They do not have direct access to their actual glucose index as reported by a physical glucometer but sense its effects (e.g., hunger) through interoceptive mechanisms involving the hypothalamus (Marty et al., 2007). Feeding is the action directed at restoring preferred glycaemic levels.

While active inference proposes a computational mechanism by which behaviour emerges from the drive to reach preferred states, hedonic evaluation provides organisms with information about the expected (wanting) and resulting (liking) values of behaviours to guide future actions (Dayan, 2022). Both the will and action of approaching preferences are positively signalled (e.g., the desire and pleasure of eating) whereas departing from preferences is negatively signalled (e.g., the displeasure of being hungry). In this proposal, we integrate hedonic evaluation and active inference to derive a general theory of how hedonic value motivates behaviour, and illustrate how the proposed framework, *hedonic foraging*, enables predicting behaviour. In the remainder of this section, we establish correspondences between the main components of hedonic evaluation and active inference, which form the foundation of our proposal, ultimately leading to hedonic foraging as a unified framework for understanding behaviour.

2.1. Hedonic impact (core liking) as variational free energy

From an active inference perspective, organisms seek states that best match their preferences (Friston, 2010). When an organism is closer to its preferred (expected) states, it experiences less surprise and hence reduced variational free energy (VFE). From the perspective of hedonic evaluation, these states are rewarding (Berridge & Kringelbach, 2013; Pessiglione & Lebreton, 2015).

We propose that VFE can thus be considered as a quantitative measure of hedonic impact or core *liking* (**Box 1**). This step constitutes our proposal's first and central component to build a framework that casts hedonic evaluation onto active inference.

As an illustration, consider a predator that expects to be in the adaptive state of satiation. When the predator is hungry, this discrepancy leads to prediction errors in its brain. This mismatch between the actual (hunger) and desired observation (satiety) is associated with high free energy. As the animal eats, this mismatch decreases, reducing VFE and eliciting a rewarding sensation. Similarly, when a person listens to Western classical music, they may experience pleasure when a perfect cadence resolves to the tonic. This fulfilment of expectations after the uncertainty introduced by the dominant chord corresponds to a state of lower VFE and a pleasurable listening experience (Cheung et al., 2019; Gold et al., 2019; Koelsch, Vuust, et al., 2019). In both examples, VFE provides a measure for the hedonic impact (liking) of the chosen behaviour. Such a quantitative reference is crucial to compare and motivate behaviour.

Box 1. Formalisation of the active inference components of hedonic foraging

Variational free energy

Active inference builds on the Bayesian brain hypothesis, which posits that the brain has an internal generative model of the world. This generative model encodes how outcomes arise from hidden causes in the environment, providing a stochastic mapping between a representation of the states of the world s and observations o. Importantly, the causal structure of the world—the generative process that describes how o is caused by the true state of the world $s*$—which may be hidden, i.e., not directly accessible to the organism (**Illustration 1a**). According to the Bayesian brain hypothesis, the organism infers the likely causes of observations by inverting its generative model using Bayes' rule $P(s|o) = P(s)P(o|s)/P(o)$.

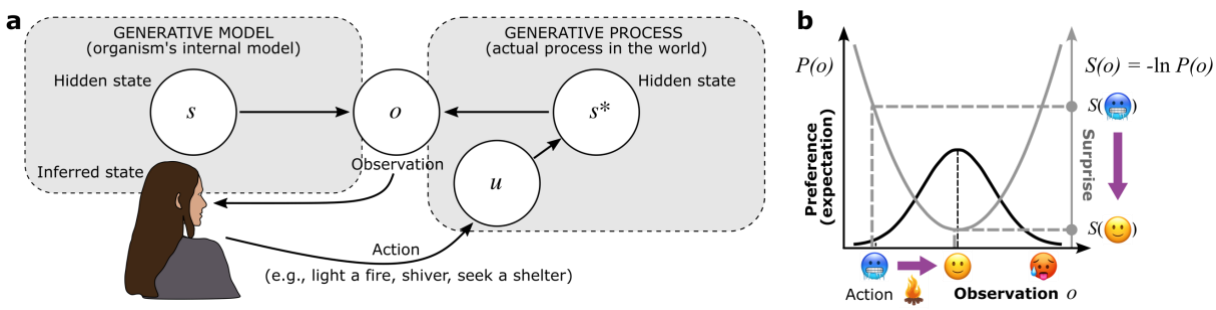

Illustration 1. a. Formalism in active inference (adapted from (Parr et al., 2022). **b.** Illustration of the correspondence between preference and surprise, with preference here driven by the homeostatic reference temperature in endotherms.

Active inference extends the Bayesian brain hypothesis with three crucial additions: First, the brain does not passively infer the causes of sensory input; it goes beyond perception to include action by actively inferring the best actions to take based on its internal model and sensory data. Second, and centrally, this inference process encompassing perception and action is governed by the imperative to minimise the deviation between actual sensory observations and preferred (i.e., expected) observations. This deviation is called surprise and quantified using information theoretic surprise: $\mathcal{S}(o) = -lnP(o)$, where P is the prior preference distribution. It is also called *prediction error* as it reflects the organism s drive to predict and sense what it prefers (**Illustration 1b**). Minimising surprise directly is often intractable, however, as the process requires inverting the brain's generative model and hence computing the posterior distribution of all possible hidden states given sensory observations. Third, to address this, the brain employs a proxy for surprise, *free energy*, enabling a computationally tractable alternative for minimising surprise. This is done by substituting $P(s|o)$ by an approximate posterior $Q(s)$, and surprise by variational free energy (VFE) F, which can be expressed as the functional

$$F[Q,o] = D_{KL}\big(Q(s) \parallel P(s|o)\big) - lnP(o) \qquad \text{(Eq 1)},$$

where the first term, the Kullback-Leibler (KL) divergence, measures the dissimilarity between the approximate distribution $Q(s)$ and the true posterior $P(s|o)$. The inference process amounts to minimising VFE. As the KL divergence is always non-negative, **Equation 1** shows that VFE is an upper bound to surprise and that there are two ways of minimising VFE: Minimising the first term, namely getting Q closer to the true posterior over states, corresponds to perceptual inference. Alternatively, the second term can be minimised by acting to change the sensory observation o. Together, active inference is the minimisation of VFE through perception and action with the aim for the organism to stay within or reach preferred states, as defined by its generative model.

Expected free energy

Planning a sequence of actions involves representing the expected future observations they will produce. In active inference, this process is guided by expected free energy (EFE), a forward-looking extension of VFE that predicts the future free energy an agent anticipates experiencing under a given sequence of actions (called a policy). EFE decomposes into terms capturing epistemic value (information gain) and pragmatic value (goal-directed outcomes) as

$$\mathrm{G}_{\pi} = \sum_{\tau>0} \underbrace{\underbrace{\mathbb{E}_{\tilde{Q}}[ln\,Q\,(A) - ln\,Q\,(A \mid s_\tau, o_\tau, \pi)\,]}_{\text{Model parameter exploration}} + \underbrace{\mathbb{E}_{\tilde{Q}}[ln\,Q\,(s_\tau|\pi) - ln\,Q\,(s_\tau|o_t,\pi)]}_{\text{Hidden state exploration}}}_{\text{Exploration}} - \underbrace{\underbrace{\mathbb{E}_{\tilde{Q}}[ln\,P\,(o_\tau)]}_{\text{Realising preferences}}}_{\text{Exploitation}} \quad \text{(Eq 2)},$$

where G_π is the expected free energy G of policy π (see Schwartenbeck et al., 2019, for details). In a discrete setting, behaviour is guided by selecting policies with the greatest potential of minimising variational free energy in the future among a finite set of options. This ensures the organism to integrate effectively exploration and exploitation to reduce VFE given current knowledge.

Habit

In active inference, habit represents the influence of past experience when selecting a policy. It is simply represented as a distribution of probability $E = (E_\pi)$ over all possible policies.

2.2. Incentive salience (core wanting) as expected free energy

While a reference measure for hedonic impact is necessary to understand motivated behaviour, it is not sufficient to predict it and thus explain how hedonic evaluation motivates behaviour. The reason is simple: VFE is retrospective; it evaluates incoming observations *a posteriori* and thus refers to past and present states. However, for planning and decision-making, an organism needs to infer the most effective future action or sequence of actions (hereafter referred to as *policies*) that generate observations yielding the most positive hedonic value, which are signalled by the lowest VFE. This cannot be done by simply considering past and current observations because actions must be chosen to bring about preferred observations in the future. To infer optimal policies, the organism must predict sequences of future states and observations and, for each possible policy, estimate its resulting VFE, that is, compute its *expected free energy* (EFE) (**Box 1**). The prospective nature of EFE enables organisms to evaluate and select policies based on their impact on future states. EFE extends VFE to account for planning, decision-making, exploration and curiosity, making it central to active inference as a theory of goal-directed behaviour (Friston et al., 2015; K. J. Friston et al., 2017; Schwartenbeck et al., 2019).

We propose EFE as a quantitative measure of incentive salience or core *wanting*. This constitutes the second component of our proposal (**Box 1**). EFE consists of two key elements: *pragmatic value* and *epistemic value* (**Box 1**, Eq 2).

Minimising the pragmatic term of EFE (Eq 2 in **Box 2**) maximises the pragmatic value, that is, the expected utility or reward associated with a future state consequence of the organism's policy. This aspect of EFE evaluates future outcomes (i.e., observations in the future resulting from policies) according to the extent to which they would fulfil the organism's goal—i.e., how close they would bring the organism to its preference. By maximising pragmatic value, organisms avoid adverse outcomes and maximise immediate reward, therefore driving *exploitative* behaviour (K. Friston et al., 2017; Schwartenbeck et al., 2019). For example, the anticipation of eating a prey enhances the likelihood of hunting it for a predator. Similarly, the anticipation of enjoying live music motivates people to go to a concert. In other words, hedonic evaluation tells us that organisms want what they believe they will like, and active inference reveals the causal relationship underlying their choice of the policy expected to best achieve their desired outcome.

Minimising the epistemic term of expected free energy (Eq 2 in **Box 2**) maximises epistemic value, which represents the expected information gain of each policy. This drives *explorative* behaviour by favouring policies that maximise knowledge about the world, resolving uncertainties about the environment and minimising prediction error. In our proposal, the epistemic component relates to incentive salience by driving the pursuit of information that reduces uncertainty about the environment, even in the absence of immediate reward (Schwartenbeck et al., 2019). For example, a predator can explore new territories to improve access to higher-quality or more abundant prey, just as humans might seek out new artists or venues to enhance future musical experiences.

In summary, minimising EFE embodies incentive motivation by increasing the likelihood of choosing policies that lead to outcomes with more positive hedonic value, both in the short and long term. This minimisation can be naturally cast onto core *wanting* (Berridge & Dayan, 2021; Berridge & Kringelbach, 2015), capturing the dual motivational pull of securing rewarding outcomes and fostering learning.

Box 2. Hedonic foraging in discrete time using POMDPs

For simplicity, we formalise hedonic foraging in discrete time[1] using a partially observable Markov decision process (POMDP, henceforth; (Malekzadeh & Plataniotis, 2024; Spaan, 2012) (**Illustration 2**), a formal structure widely used in active inference to model organisms (or, more generally, agents) operating under environmental uncertainty and partial observability (Smith et al., 2022). This framework captures the interaction between an organism's *generative model* (its internal representation of the world) and the *generative process* (the actual dynamics of the environment) (**Box 1**). Under the free-energy principle, decision-making unfolds as the organism updates its beliefs about *hidden states* (s_t), which evolve over discrete time steps, with the objective of minimising EFE. Actions are selected based on policies (series of actions) that maximise the probability of achieving preferred *outcomes*, or *observations* (o_t). These observations can include exteroceptive, proprioceptive or interoceptive sensory input, as well as cognitive representations.

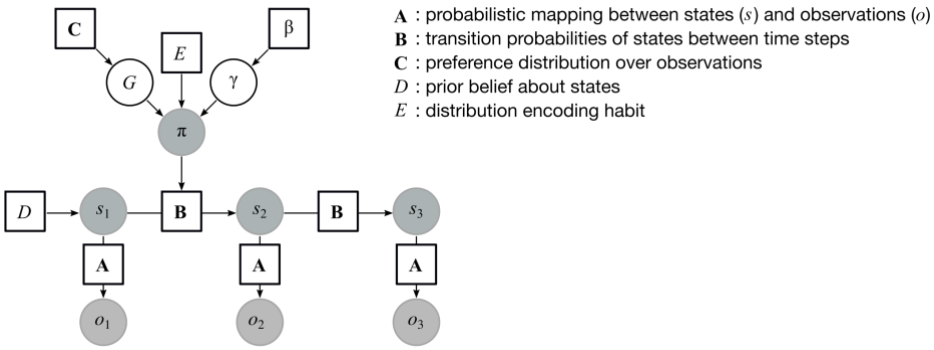

Illustration 2. Basic formalism of partially observable Markov decision processes (POMDPs).

By minimising EFE, the organism achieves preferred outcomes (pragmatic value) and continuously refines its generative model, improving predictions while selecting actions that reduce uncertainty (epistemic value). The dynamics of a POMDP—and each scenario in **Section 3**—are specified in terms of a set of matrices or tensors ($\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$) and vectors (D, E) representing the alignment between the generative model and the generative process. When they align, the organism's generative model exactly captures the generative process leading to its observations—although uncertainty about initial states and stochasticity between states and observations may still exist. When they do not align, the organism can learn the generative process through repeated observations in a process akin to Hebbian learning (e.g. see preference learning in **Figure 3b**).

Hedonic foraging is thus described by the iteration of two steps: First, the organism selects a policy according to the prior distribution over policies $\sigma(lnE - \gamma G)$, which combines habit and EFE to represent the "value" of all the policies as a discrete probability distribution—here, σ is a softmax function $\sigma(x)_\pi = e^{-\alpha x_\pi} / \sum_k e^{-\alpha x_k}$ for $x = (x_\pi)$, α is a temperature parameter that controls action precision (as high values increase the differences between entries), and γ is the EFE precision parameter that encodes the organism's confidence in the ability of EFE to predict VFE. The second step is given by the posterior probability distribution over policies $\sigma(lnE - \gamma G - F)$, which combines habit, EFE and the VFE resulting from the new observations.

When these two steps are completed, two key processes take place: First, the EFE precision parameter is updated based on the agreement between the prior and the posterior, with an increase in precision (γ) when they align (Sales et al., 2019; Smith et al., 2022). Second, all learning parameters in the model are updated. This learning process resembles Hebbian learning, as it relies on coincidence detection. For instance, if the habit component is subject to learning, its weight in the habit probability distribution E_{π_0} increases every time a specific policy π_0 leads to a positive outcome. Consequently, this policy becomes more likely to be selected in the future.

Importantly, the influence of wanting in guiding behaviour is modulated by its capacity to lead to liking after the consequences of a selected policy have been experienced and evaluated. As implemented in active inference, the contribution of EFE in defining the prior is weighed by the confidence the organism places in its ability to predict rewarding behaviour, that is, in minimising VFE. To this end, the prior is computed using the distribution over policies $lnE_\pi - \gamma G_\pi$, where γ is the precision parameter that encodes the organism's confidence in EFE estimates (**Box 1**). Every time new observations are received, the precision parameter is updated. If the posterior carried by VFE aligns with EFE—i.e., the policy with lower EFE before any observation was available corresponds to the one that actually leads to a reduction in VFE—, γ increases. The organism is then more confident in the beliefs provided by G_π, and the impact of EFE (wanting) in guiding behaviour grows. The opposite happens when the VFE of the resulting observations is inconsistent with the estimate provided by EFE, in which case the confidence of the organism in its beliefs brought by G_π decreases. Dopamine release signals such confidence in policy selection (Bogacz, 2020; FitzGerald et al., 2015; Friston et al., 2012) by ramping according to wanting in the appetitive phase (**Box 1**).

[1]Hedonic foraging can also be expressed in continuous time using the continuous-time formulation of active inference (Parr et al., 2022).

2.3. Learned hedonic values as habits

In active inference, learning occurs at multiple time scales (Friston et al., 2016). A particularly interesting form of learning emerges from repeated behaviour, which may—although does not necessarily—result from consistently selecting policies that have proven beneficial in the past (Thorndike, 1911). *Habits* emerge from a history of repeated behaviours, reflecting the role of past actions and their associated hedonic values in motivating future behaviour (Seger & Spiering, 2011) (**Box 1**).

As such, habit constitutes an essential drive, capturing long-term associations between expectations and outcomes to complete the map of interrelations between hedonic values and the constructs of active inference. While EFE drives goal-directed behaviour, characterised by adaptive flexibility to environmental contingencies (Dayan, 2009), habit accounts for recurrent behaviour defined by past reinforcements but possibly unrelated to current hedonic values; being thus automatic, computationally efficient and inflexible (Dolan & Dayan, 2013).

Habit formation provides adaptive advantages, as habits may maximise efficiency in policy selection for recurrent scenarios when there is little evidence from data, as they minimise energetic costs and risks when the available options are invariable and no further exploration is required or desired (Dayan, 2009; Dolan & Dayan, 2013). For example, the habit of visiting a particular location at a particular time of day or year when prey is abundant can be greatly beneficial for a predator as it does not have to spend energy and take risks to find it. Similarly, the habit of attending concerts at a known venue may increase the odds to enjoy a pleasurable musical experience, as the listener has learned that previous experiences have been consistently pleasurable.

2.4. Motivated behaviour as hedonic foraging

The central idea of this proposal is that the mechanisms of motivated behaviour can be characterised as an interaction between habit, EFE minimisation (wanting) and VFE minimisation (liking). First, policy selection depends on the combination of habits (E, reflecting a history of previous policy selections and their resulting hedonic impacts: liking or disliking) and expected free energy (G, corresponding to wanting for available policies). Then, the selected policy leads to new observations, whose VFE (F, corresponding to liking or disliking the resulting outcomes) is used to evaluate the effectiveness of the policy. In this way, the hedonic components of habit and wanting define the prior expectations over policies, while their combination with liking determines the posterior for the selected policy, which in turn will inform subsequent priors. We refer to **Box 2** for a detailed formalisation of hedonic foraging.

Hedonic foraging can thus be understood as navigating a free energy landscape (**Figure 1**), using the probabilities of selecting actions provided by habits and EFE as a compass, with continual adjustments informed by the VFE resulting from each policy. At each step, the organism chooses a policy that optimises, with some stochastic variability, a balance between habit and wanting. After each policy, the organism evaluates the VFE based on the resulting observations, which is reflected in a degree of liking or disliking. This hedonic evaluation then updates the active inference process, modulating the organism's habits and wanting and, more generally, refining the generative model that guides its actions. Behaviour thus unfolds iteratively, with each iteration consisting of these two interdependent stages (**Figure 1**).

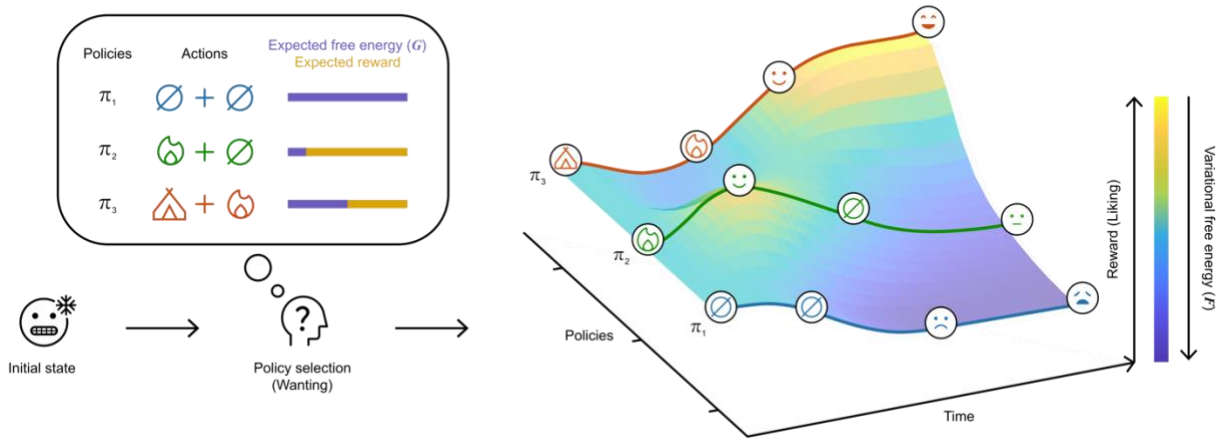

Figure 1. Hedonic foraging schematic: Free energy and reward landscape in thermoregulation
As endotherm organisms, humans must keep their body temperature within a specific range. Cold, in particular, poses a significant threat, as prolonged exposure is life-threatening. The unpleasant sensation of cold increases the motivation (wanting) to restore body temperature. Various policies differ in their effectiveness, effort and time required to achieve this goal. For example, a hiker suddenly confronted with an unpleasant cold sensation but not taking any action (π_1) would likely remain cold, further intensifying the unpleasant sensation (as VFE increases). However, she may take action to reach a more beneficial state. For instance, she could light a fire (π_2) to raise her body temperature quickly, providing an immediate but transient reward (as VFE decreases), since her body temperature drops again when the fire extinguishes, again eliciting the unpleasant sensation of cold (as VFE increases). Alternatively, she may consider the more elaborated, time-wise and energy-wise costly, deeper policy (π_3) of building a shelter (initially increasing VFE and hence discomfort) before lighting a fire, ultimately leading to a greater and more lasting reward (as VFE decreases). Overall, the first policy (π_1) would be disliked while the second (π_2) would be mildly liked and the third (π_3) would be strongly liked. By monitoring policies and their expected and experienced hedonic values at each time step, organisms can learn embodied associations between (series of) actions and outcomes, using this knowledge to guide future decisions and behaviours.

Hedonic foraging casts behaviour as a balance between habits and EFE with the aim to minimise VFE, leading to various scenarios. In simpler cases, one component may dominate when the others are indifferent to the outcome of a new policy: e.g., a habit may drive behaviour when all options are equally appealing and there is no incentive salience for learning—the distribution of expectations of EFE is flat. In more complex cases, a chosen policy may reflect a trade-off between competing components: e.g., when habitual behaviour conflicts with the drive for epistemic exploration. In the next section, we illustrate the scope of our proposal across a range of scenarios.

3. A continuum of behaviours through hedonic foraging

In this section, we demonstrate that a continuum of behaviours can be captured within the single, unifying framework of hedonic foraging. To achieve this, we simulate behaviour across four illustrative scenarios. The model's architecture and the computational mechanisms that govern behaviour are the same across scenarios. This consistency is the basis for describing the model as accounting for a continuum of behaviours. Only the parameters of the generative models and the generative processes are specific to each behavioural scenario.

The following set of mappings defines the generative models characterising all scenarios: The first matrix (or tensor, $\boldsymbol{A}$ in **Box 2**) maps hidden states to observable outcomes. This mapping is probabilistic, providing the likelihood of a particular observation given the true state of the world. For example (**Figure 1**, **Illustration 1b** in **Box 1**), a state of hypothermia leads to the sensation of cold most of the time but not always. The second mapping ($\boldsymbol{B}$ in **Box 2**) describes the probabilistic transitions between hidden states from the present to the next time step, depending on each action taken. In the example in **Figure 1**, not moving keeps the agent in the state of hypothermia, whereas looking for shelter entails transitioning to the state of normothermia with a probability of 0.8. Finally, the preference matrix ($\boldsymbol{C}$ in **Box 2**) is central as it encodes the agent's desired outcomes as probability distributions, fundamentally shaping the VFE landscape and, consequently, the EFE. In this example, the organism has a strong (adaptive) preference for staying in the reference (homeostatic) state of normothermia, which translates into positive hedonic values (desire, reward) for such reference state and negative hedonic values (avoidance, punishment) for the states of hypo- and hyperthermia.

In the POMDP-based scenarios presented below, the behaviour of the simulated organism emerges dynamically from the interaction between the three components of hedonic foraging. Hedonic foraging is characterised as the stochastic evolution of the organism's states and observations, driven by the combined influence of all POMDP factors—including observation mappings, state transitions, habits and preferences—, each constraining the organism's choices in response to its environment. The code implementation and detailed description of each scenario are publicly accessible at https://osf.io/v39mb/?view_only=852e26bf70da4f7aa919ea41fb20fcd5.

Scenario 1: Allostasis

In this scenario, we simulate how hedonic foraging accounts for the fundamental process of allostasis in a simple setting: glycoregulation (López-Gambero et al., 2019; Roh et al., 2016). An organism can be in one of three physiological states: hypoglycaemia, normoglycaemia or hyperglycaemia. It estimates these states through interoceptive observations: the *sensations* of hunger, satiety and indigestion, respectively. The mapping between hidden states and sensations can be noisy, meaning that a true state does not always align perfectly with its usual sensation (e.g., hypoglycaemia can be perceived as satiety). The organism iteratively chooses between two actions: eating increases the likelihood of transitioning to a higher glycaemic state (from hypoglycaemia to normoglycaemia and from normoglycaemia to hyperglycaemia) whereas resting increases the likelihood of transitioning to a lower glycaemic state (from hyperglycaemia to normoglycaemia and from normoglycaemia to hypoglycaemia). Ultimately, the organism strongly prefers satiety (normoglycaemia). As a result of hedonic foraging, the organism's feeding cycle emerges from the imperative to reach the preferred state and avoid surprise (**Figure 2**).

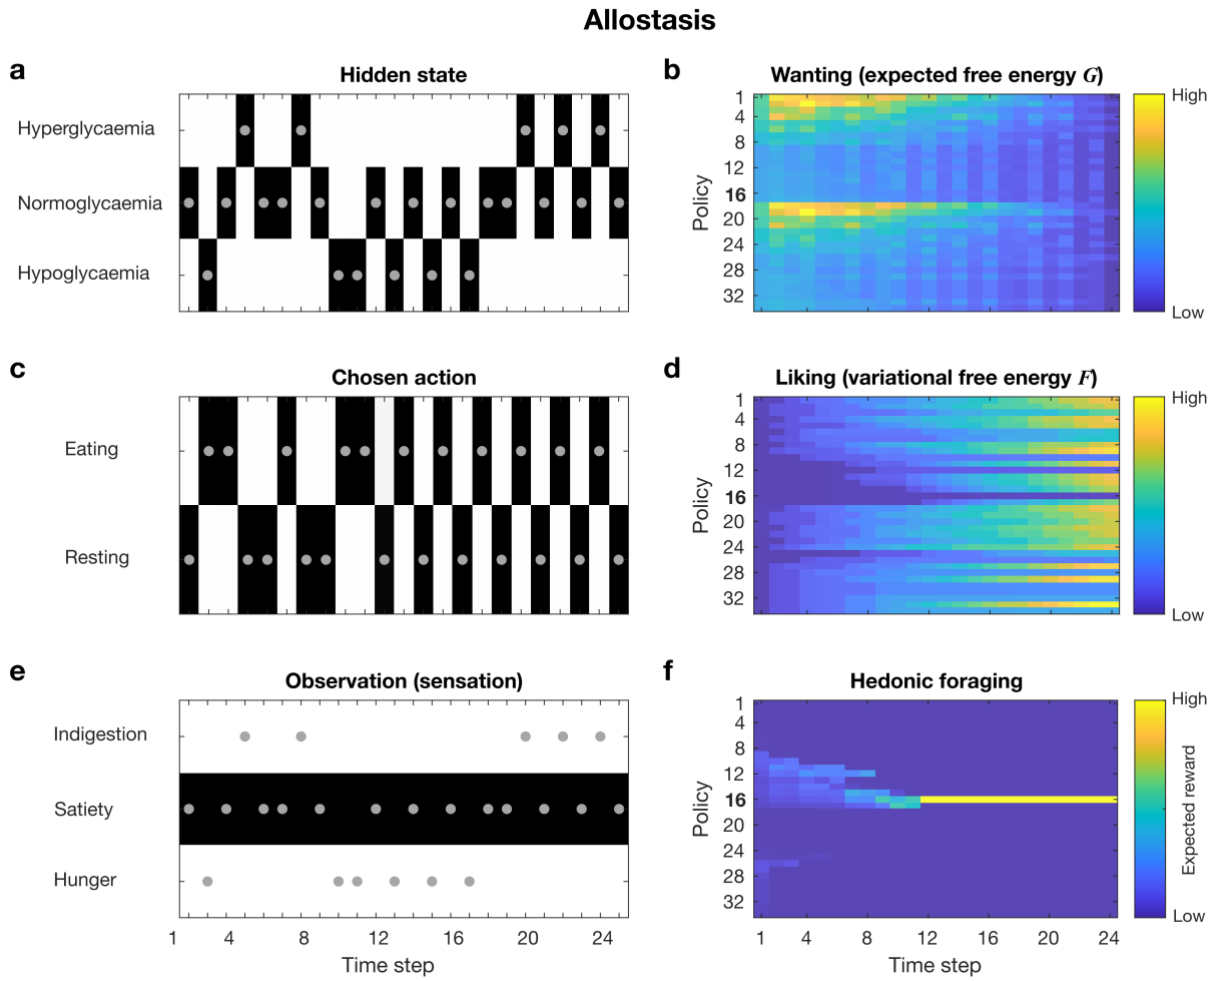

Figure 2. Allostasis (glycoregulation)

In this instance of glycemia regulation, the organism starts in normoglycemia (t_1). When its glycemic level decays (t_2), it enters a hypoglycaemic state and hence feels hungry, so it starts eating (t_3) to reach a normoglycemic state and feel satiated. Eating in normoglycemia (t_4) leads to hyperglycaemia, with the accompanying sensation of indigestion. Resting (t_5) leads to progressive decay in glycemic level until reaching the preferred normoglycemic state, which progressively decays into a hypoglycaemic state, triggering eating again (t_6). The left panels show how a feeding cycle emerges from the goal of maintaining normoglycemia as long as possible. Here, dots represent the hidden state (**a**), the chosen action (**b**) and the resulting observation (**c**) at each time step. The process is not deterministic, so the organism can, for example, keep eating even when satiated. The policy choices are driven by hedonic foraging, with the greatest reward resulting from eating when feeling hungry. The right panels show the dynamics of wanting (G, **b**), liking (F, **d**) and hedonic foraging (**f**), with lower expected free energy corresponding to higher expected reward and variational free energy corresponding to actual reward: the choices converge into a stable policy (π_{16}) yielding maximum global hedonic value.

Scenario 2: Statistical learning and preference formation

Statistical learning is the process by which an organism implicitly learns regularities in sensory input, shaping future expectations (Turk-Browne, 2012). When an input aligns well with learned patterns, this match triggers a positive affective response, as the sensory-cognitive systems are particularly suited to processing familiar features efficiently (Bornstein & Dagostino, 1992; Reber et al., 1998). Statistical learning thus subserves *mere exposure* effects, which refer to a tendency to develop liking for stimuli as a function of familiarity (Zajonc, 1968). These effects are shared with other animals and have been extensively demonstrated in humans to influence the hedonic evaluation of a broad range of stimuli, from food (Birch & Marlin, 1982; Pliner, 1982) to music (Pereira et al., 2011) and artworks (Leder, 2003; Reber et al., 2004).

This scenario exposes a simple link between statistical learning and preference formation. Here, an organism freely chooses among three equally preferred locations. Two contain stimulus A while the other contains a different but equally rewarding stimulus B. That is, the organism has no initial preference for either stimulus or location over the other. Without any incentive to favour a specific location, stimulus A is twice as likely to be encountered. According to the mere exposure effect, a greater exposure reinforces preference for stimulus A over B. As the organism progressively learns the frequency of each stimulus and where to encounter them, the pursuit of reward creates a feedback loop in which frequent visits to locations with the preferred stimulus strengthen the preference for that stimulus, further increasing the likelihood of returning. This behaviour ultimately leads to establishing a habit of visiting the locations typically containing the preferred stimulus.

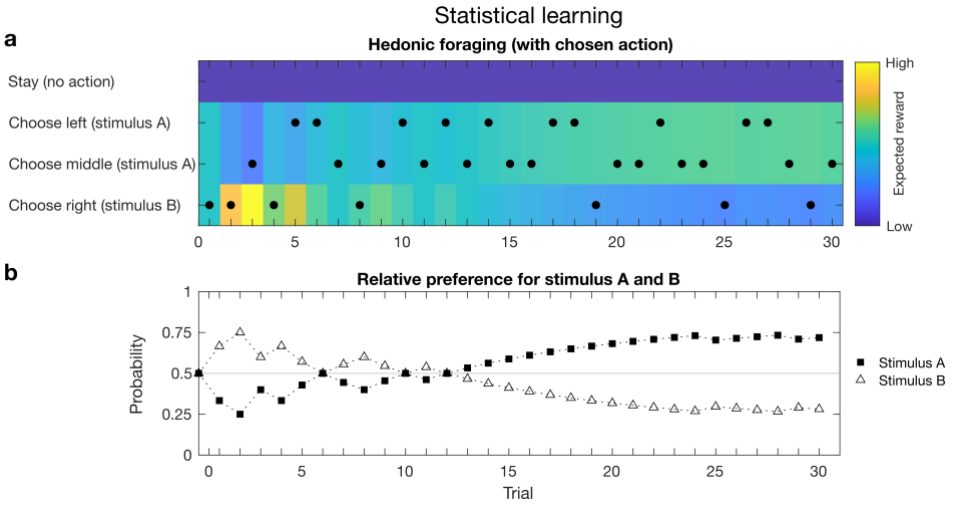

Figure 3. Statistical learning and formation of familiarity-driven preferences
In this example, the organism initially has no preference between stimulus A and B, leading to random choices. Over time, it learns the frequency and locations of each stimulus: A appears in 2/3 of locations while B does in 1/3. Even if the agent first encounters B (t_0), it quickly infers that A is more frequent and shifts its preference accordingly, increasing visits to A-associated locations ($t_{>1}$). Hedonic foraging drives behaviour by selecting policies (dots in **a**) that maximise reward according to updated beliefs about reward probability (trajectories for each stimulus in **b**).

Scenario 3: Bias and curiosity

Hedonic evaluation is context-dependent (Berridge, 1991; Lim, 2011; Skov, 2023). Through statistical learning, organisms identify regularities in spatial and temporal co-occurrences (Turk-Browne, 2012) and develop preferences for familiar patterns, as shown in **Scenario 2**. Consequently, preferences for specific stimuli can be *biased* by their co-occurrence with other stimuli or contextual features with enhanced evaluative salience (Lopez-Persem et al., 2016). For instance, the *white cube effect* refers to a tendency to prefer artworks displayed in museums over those encountered in public spaces (Brieber et al., 2015; Gartus & Leder, 2014). However, *curiosity* (Gottlieb & Oudeyer, 2018; Kidd & Hayden, 2015) can counteract such biases. This mitigating effect of curiosity is accounted for in hedonic foraging, which incorporates curiosity via the epistemic term of expected free energy—fostering wanting to explore unknown aspects of the environment to reduce uncertainty (Schwartenbeck et al., 2019).

In this scenario, an organism (e.g., a museum visitor) can choose between two contexts leading to equally pleasurable stimuli (e.g., a particular artwork). The organism prefers one of the contexts (e.g., a museum over a metro station) because it has statistically, culturally or socially learned that the stimuli associated with this context are usually more rewarding than those typically present in the other context. The incentive motivation (G in **Box 1**) or habit ($\boldsymbol{E}$ in **Box 1** and **Box 2**) to visit the preferred context more often bias the hedonic impact (F in **Box 1**) of stimuli in that context, thus enhancing their preference ($\boldsymbol{C}$ in **Box 2**) or reinforcing the habit of visiting that context ($\boldsymbol{E}$ in **Box 1** and **Box 2**). **Figure 4** shows contrasting levels of curiosity and demonstrates that more curious organisms visit the less-preferred context more often (**Figure 4a** as opposed to **Figure 4b**), resulting in bias reduction and even bias reversal, as these individuals have the opportunity to evaluate the stimuli in their own right.

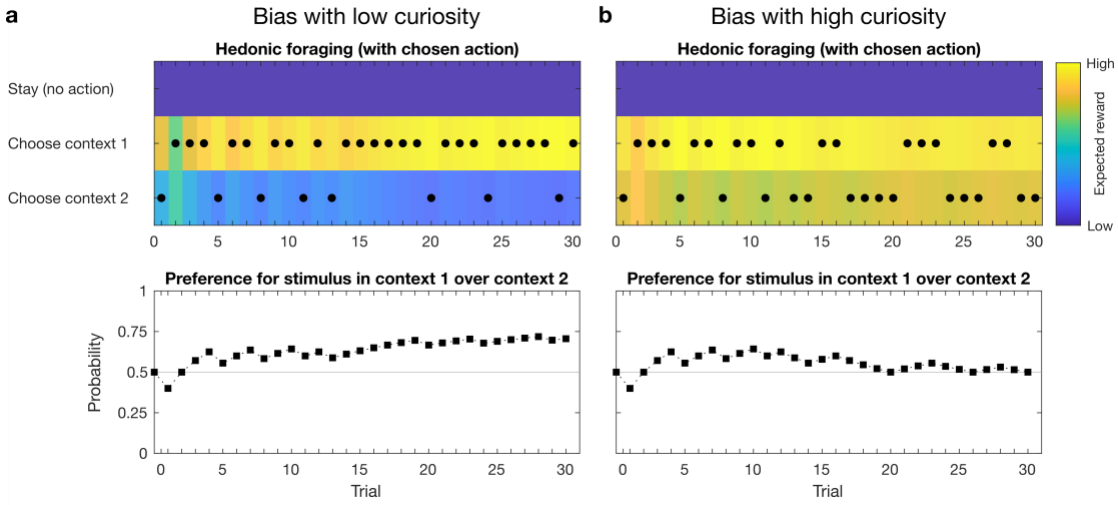

Figure 4. Interplay between bias and curiosity
In this example, contexts 1 and 2 contain the same stimulus. The organism has an initial bias (preference) for context 1, even if it starts visiting context 2 due to the stochasticity of policy choice (t_0). Hedonic foraging guides behaviour, with policy selection (top panels) maximising reward based on updated beliefs about reward probability (bottom panels). **a.** When the curiosity to explore other contexts is low, the agent visits context 1 more frequently (top panel) because it believes this context yields greater reward, which in turn reinforces this belief (bottom panel). **b.** By contrast, higher curiosity counteracts this bias, as exploration allows the agent to experience the same stimuli in both contexts and evaluate them in their own right. Increased exploration reveals that stimulus value is context-independent, leading to more balanced use of both contexts and the vanishing of prior biases.

Scenario 4: Insight, *Erlebnis* and “Aha” experiences

Ambiguous stimuli elicit unstable percepts (Klink et al., 2012). From the predictive coding perspective within active inference, uncertainty generates prediction error (Cheung et al., 2019; Gold & Shadlen, 2007; Koelsch, Bashevkin, et al., 2019; Ross & Hansen, 2016). Organisms seek to resolve such uncertainty and thus reduce prediction error, deriving reward from the process. When this understanding occurs suddenly, with a feeling of ease and certainty (Topolinski & Reber, 2010), usually requiring only a few observations and favoured by curious behaviour (K. J. Friston et al., 2017), it is called *insight*, *Erlebnis* or “Aha” experience. Gaining insight or understanding can occur exogenously, through external cues or explicit information, or endogenously, through internal reasoning or exploration—each with a distinct probability of success. Organisms vary in their preferences for seeking insight either exogenously or endogenously and in their ability to gain insight with or without external cues—see (Kounios & Beeman, 2014) for a review of the cognitive neuroscience of insight in humans and (Shettleworth, 2012) for a discussion of the phenomenon in other animals with a brain.

This scenario shows how individual ability and relative preference for exogenous or endogenous insight guides an organism’s behaviour. It can be situated, for example, in a museum where visitors choose whether to consult the labels accompanying the artworks (**Figure 5**).

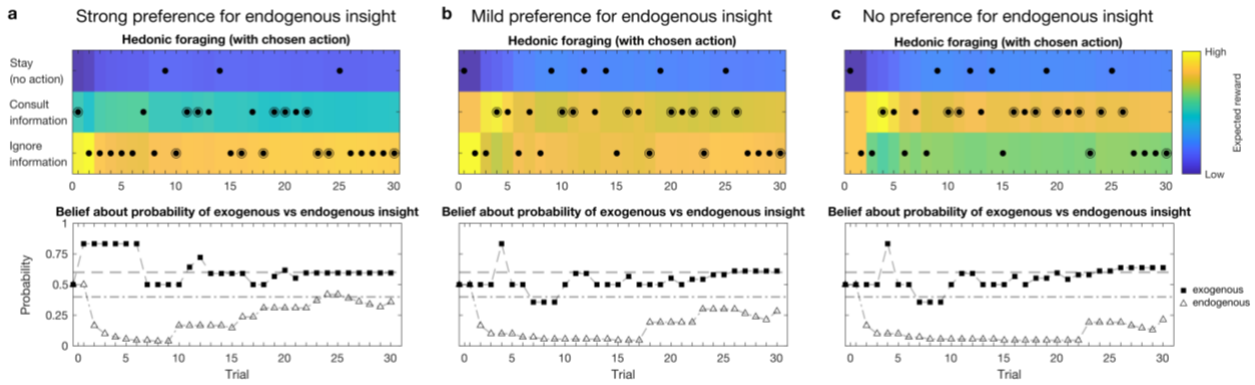

Figure 5. Preference for endogenous vs exogenous insight
An organism is presented with an ambiguous stimulus about which it seeks to gain insight. This can be achieved through two strategies: exogenously, by consulting external information, which increases the probability of insight, or endogenously, relying on internal processing with a lower probability of insight. Organisms also exhibit individual variability in preference for endogenous over exogenous insight, as well as distinct ability (probability) to gain insight (circled dots) through each modality (0.6 for exogenous and 0.4 for endogenous across panels). Hedonic foraging captures how such traits determine individual behaviour (dots represent chosen actions in the upper panels, while the lower panels depict the evolution of beliefs about the relative probabilities of gaining insight). **a.** When the preference for endogenous insight is significantly stronger than the preference for exogenous insight, the expected reward outweighs the risk of not getting insight, leading the organism to disregard external information and trust its own capability more often. **b.** When the preference for endogenous insight is not strong enough to compensate for the risk of not getting insight, both strategies are more balanced. **c.** If there is no preference for endogenous insight, the organism prioritises gaining insight, relying more frequently on exogenous facilitation.

4. Contextualising hedonic foraging

In this section, we review theoretical frameworks and computational models that explore the relationships between hedonic evaluation—along with the related concepts of emotional valence and affective value—and motivated behaviour. We categorise these approaches into those grounded in active inference and those based on alternative theoretical paradigms. For each approach, we highlight its core principles, summarise its key findings and evaluate how it relates to hedonic foraging.

4.1. Emotion, affect and epistemic curiosity

Theories and models built on active inference that aim to establish principled links between hedonic value and behaviour have primarily focused on the subjective emotional impact of experiences—specifically, on *emotional valence,* the positive or negative quality of emotion. Joffily and Coricelli (Joffily & Coricelli, 2013) proposed an association of emotional valence with the rate of change of VFE. As in hedonic foraging, their framework introduces a meta-construction in that it applies to any active inference process independently of the underlying specific generative model. In their proposal, emotional valence emerges from the dynamics of VFE, tracked through its first and second derivatives. These dynamics are directly tied to the organism's adaptation to its environment. In contrast to hedonic foraging, the association between emotional valence and free energy in Joffily and Coricelli's proposal was not extended to action, limiting its ability to explain or predict motivated behaviour. The reason is that this framework predated the development of EFE in active inference, which provides a formal mechanism to link affective states with action selection (Friston et al., 2015).

Emotional valence and adaptive behaviour have also been connected through active inference via the theory of *constructed emotion*, which posits that emotions emerge as holistic, context-dependent brain-body phenomena (Barrett, 2017). This theory highlights the role of predictive processes in constructing emotional experiences, with the brain dynamically modelling internal and external states to regulate the body (allostasis) and interpret sensory inputs. As part of this process, emotional valence reflects the brain s predictions about bodily and environmental changes. Hesp and colleagues (Hesp et al., 2021) extended this framework by introducing *affective charge*, a computational perspective that incorporates meta-cognition, explicitly linking confidence in action selection to affective valence and subjective fitness through second-order beliefs—i.e., beliefs about beliefs. Affective charge can therefore be understood as a (signed) reward prediction error, where positive prediction errors increase the agent's confidence in its action model and vice versa (Hesp et al., 2021). Hedonic foraging, by contrast, emphasises the direct role of habit, wanting and liking in shaping behaviour. In sum, affective charge highlights the intrinsic link between affect, implicit metacognition and action, where biological systems infer and regulate their own affective states, while hedonic foraging focuses on the primary forces driving behavioural choices. The idea of affective charge can be integrated into the hedonic foraging framework as second-order monitoring of reward maximisation, i.e., of free energy minimisation.

The relationship between uncertainty reduction and positive hedonic value—or reward—has also been explored as a mechanism for understanding art appreciation. These studies are framed within the theory of predictive coding, which serves as a key foundational component of the broader framework of active inference. Predictive coding suggests that the brain seeks to make sense of sensory input by inferring the hidden states of the environment through Bayesian

inference. Under this framework, predictions from a generative model are compared to incoming sensory information, with the mismatch—prediction error—being reminiscent of VFE in active inference. Applying this to art appreciation, Van de Cruys (Van de Cruys, 2017) proposed that transitioning from a state of higher uncertainty to one of lower uncertainty—thus reducing prediction error—and positive rates of prediction error reduction elicit positive affective experiences. According to this view, the intense pleasure of "Aha" or *Erlebnis* moments arises from such a resolution of uncertainty. This perspective aligns with our proposal, where liking is proportional to the reduction in VFE. However, (Van de Cruys, 2017)'s framework did not address why organisms engage in behaviours that can initially increase prediction error, as seen in exploratory or epistemic actions.

Overcoming this limitation, Van de Cruys and colleagues (Van de Cruys et al., 2024) have recently proposed that engaging in behaviours that initially increase prediction error is explained by the epistemic component of curiosity, which reflects expected information gain and the expectation of resolving uncertainty. They describe this process as an *epistemic arc* comprising three interconnected stages: curiosity, epistemic action and the resulting "Aha" experiences, which they interpret as the foundation of *aesthetic* experiences. While their model draws on the process theory of active inference and acknowledges potential applications beyond art-related experiences—such as everyday sense-making processes—it remains largely conceptual and lacks a quantitative framework. Our approach addresses this gap by providing a quantitative method for characterising epistemic arcs, situating them within a continuum of experiences not restricted to specific domains or species. Our formulation builds on active inference by offering a precise characterisation of the epistemic and pragmatic motivations that drive engagement with the world. By extending these principles, our approach underscores the importance of formal models in capturing the general mechanisms underlying curiosity-driven behaviours across diverse contexts, and how hedonic evaluation monitors and drives such behaviours.

4.2. Other models of hedonic value

Several approaches unrelated to active inference have proposed links between the processes involved in hedonic evaluation and motivated behaviour. Skov (Skov, 2023) offered a conceptual model explaining how nervous systems assign hedonic value to sensory objects, leading to action. As in the current proposal, hedonic values are inherently adaptive, shaped by the organism's survival needs and ever-changing environment. Liking and disliking trigger response patterns that enable organisms to recognise and react to situations that could aid or harm their survival. Therefore, sensory valuation is not a simple reflection of an object's intrinsic properties but depends on the organism's idiosyncratic and current needs and evaluative context. For example, eating may elicit a positive hedonic value when the organism is in hypoglycaemia (as in **Scenario 1** in **Section 3**) and a negative value when it is in hyperglycaemia (López-Gambero et al., 2019; Roh et al., 2016). While Skov s model remains conceptual and does not accommodate data or make predictions, its conceptual similarities with our proposal and its detailed description of the neurobiology of hedonic evaluation make it a valuable resource for refining and enhancing descriptions based on active inference.

Our proposal shares several features with computational models of hedonic value based on reinforcement learning. Both reinforcement learning and active inference are theories of decision-making. In reinforcement learning, the objective is to determine the optimal policy—the one that maximises reward—through a value function following Bellman s equation (Sutton & Barto, 2018). This function includes components for immediate rewards and expected future rewards. In a proposal focused on hedonic value (referred to as *aesthetic* value), Brielmann and

Dayan (Brielmann & Dayan, 2022) mapped two processes known to contribute to art appreciation to these two components of reinforcement learning: processing fluency (Reber et al., 2004; Reber et al., 1998) informs immediate rewards while learning influences expected future rewards by facilitating future processing efficiency. Drawing on this parallel, the authors suggest that art experiences are processed in the brain similarly to how other rewarding experiences are handled in reinforcement learning, with hedonic value serving as a reward signal that guides future behaviour and preferences. While this proposal establishes a formal analogy with *wanting* and *liking* in hedonic evaluation, it does not account for the intrinsic motivation in learning or the pre-existing preferences that drive behaviour. This limitation is inherent to all reinforcement learning models due to their lack of an explicit formulation of epistemic (uncertainty-reducing) and pragmatic (goal-directed) imperatives, which are present in active inference (K. J. Friston et al., 2017; Pezzulo et al., 2024). In contrast with reinforcement learning, agents in active inference are endowed with a model of purposive behaviour: they act to fulfil their prior preferences. In addition, unlike value functions, which must be inferred, active inference naturally incorporates degrees of belief and uncertainty, which are central to understanding adaptive decision-making (Doya et al., 2007). In that sense, active inference offers a compelling alternative to reinforcement learning for analysing and generating motivated behaviour as it focuses on an agent's beliefs, uncertainty and intrinsic motivations rather than externally defined rewards. In so doing, active inference provides deeper insights into exploration, decision-making under uncertainty, and adaptability in dynamic environments (Friston et al., 2016; Pezzulo et al., 2024; Schwartenbeck et al., 2019).

5. Why should we expect a continuum?

Hedonic foraging integrates hedonic evaluation into a single, unified model of action and perception—active inference—to describe a comprehensive range of behaviours that transcends biological taxa. In this framework, the core constructs of active inference remain consistent, with only the modulating parameters being scenario-specific and tailored to each organism and behaviour. This is how *continuum* should be understood here: a single framework accounts for a range of observed phenomena, illustrating continuity across behavioural complexity and species endowed with an evaluation system.

But why should we expect a continuum in the first place? This enquiry is central to many disciplines, as it carries fundamental scientific and philosophical implications for understanding cognition and human nature. In particular, it addresses whether there is a computational or neurophysiological mechanism of behaviour unique to *Homo sapiens*. If such a special mechanism exists, what does the break in the continuum look like, mean and entail for cognition? If a common mechanism drives a continuum of behaviours across scenarios and taxa, even sophisticated human pursuits and experiences traditionally deemed distinctly human (e.g., the so-called *aesthetic* experiences) would obey the same principles as allostatic processes in simpler organisms such as worms. We can refine the question to resolve this conundrum as follows: Should we expect *wanting* or *liking* nutritionally valuable food (a primary reward satisfying biological needs) and *wanting* or *liking* in the conscious appreciation of an artwork (a secondary reward without direct biological relevance) to share the same computational and neurophysiological mechanisms?

Scholars in empirical aesthetics and related disciplines have traditionally insisted on a negative answer to this question, based on the widespread but controversial belief that art appreciation is disconnected from other psychological and neurological processes, as critically reviewed by Skov and Nadal (Nadal & Skov, 2024; Skov, 2019; Skov & Nadal, 2018, 2020). However, a central argument for a broadly shared basis of hedonic evaluation in motivated behaviour rests on substantial evidence that not only a wide array of pleasurable experiences—e.g., eating chocolate, having sex, listening to music—engages the same neural circuits in humans, but also that these circuits are evolutionarily conserved and widely shared with other animals (Berridge & Dayan, 2021; Berridge & Kringelbach, 2013, 2015). This supports the hypothesis of a *common currency* for hedonic evaluation (Levy & Glimcher, 2012; Montague & King-Casas, 2007), according to which first and secondary rewards are processed within a shared neural network across a range of species (Bartra et al., 2013; Sescousse et al., 2013).

From an evolutionary perspective, the function of *liking* is not to generate positive introspective experiences but to enable organisms to identify and respond to adaptive opportunities and threats (Pessiglione & Lebreton, 2015; Rosenthal, 2017). Darwin (Darwin, 1872) anticipated this idea, proposing that emotional reactions are primarily selected for their adaptive value. This view has been substantiated by modern neuroscience, which provides extensive evidence that affective processing is underpinned by a core set of mechanisms shared across phylogeny (Anderson & Adolphs, 2014; LeDoux, 2012). Over time, cognitive and cultural pleasures have co-opted the hedonic circuitry initially evolved for basic sensory pleasures (Berridge & Dayan, 2021; Nguyen et al., 2021). Core hedonic values (affective signals) appeared first, followed by subjective and conscious evaluations computed by more complex, hierarchical structures (Berridge & Kringelbach, 2015; Panksepp, 1998). It has been proposed that wanting evolved initially to pursue innate rewards, such as nutrients and mates, while liking likely emerged later

to evaluate behaviour and inform subsequent wanting (Berridge & Dayan, 2021; Winkielman & Berridge, 2003). This would explain both innate drives and exploration of novel environments, with learning resulting from the evaluation of such exploration. In these cases, reward maximisation (EFE minimisation) would have been first initiated by uninformed predictions (hypothetical VFE) and later informed by previous experiences (resulting VFE).

Crucially, hedonic foraging describes a common computational mechanism implemented through specific neurobiological processes. While the framework suggests that the principle of reward maximisation spans behaviours and species, the scenarios and the brain structures underlying hedonic evaluation differ between taxa. More complex brains enable more refined cognitive abilities and sophisticated computations, such as deeper policy evaluation and planning: for example, *C. elegans* search for food using simple exploration strategies (Calhoun et al., 2014; Madirolas et al., 2023) whereas orcas plan and coordinate collective hunts (Lopez & Lopez, 1985; Pitman & Durban, 2012). Furthermore, distinct brain structures subserve similar functions in different species (see (Jarvis et al., 2005), for a review of brain evolution in vertebrates): for instance, executive functions involved in behavioural optimisation are associated with activity in the mammalian prefrontal cortex (Chudasama, 2011; Friedman & Robbins, 2022) and the avian nidopallium caudolaterale (Güntürkün, 2005; Güntürkün & Bugnyar, 2016).

However, despite specific differences, the neural machinery underpinning hedonic evaluation is conserved across phylogeny, with biogenic amines (e.g., dopamine) playing a fundamental role in behavioural responses to rewards across various animal phyla (Barron et al., 2010). The emergence of the mesolimbic reward system predates the divergence of chordates (Yamamoto & Vernier, 2011) and there is substantial converging evidence for homology across major vertebrate lineages (O'Connell & Hofmann, 2011). Dopamine and other neurotransmitters mediate hedonic signals in sensory learning and motivate behaviour in invertebrates as well (Unoki et al., 2006). Moreover, neurogenetic studies suggest that the insect reward system has a modular structure that broadly parallels the mammalian system in terms of functional organisation, with distinct brain circuits for reward learning, wanting and liking (Perry & Barron, 2013).

Together, these functional, phylogenetic and neurophysiological similarities suggest a unified framework for understanding how hedonic evaluation guides behaviour across scenarios and taxa. As reviewed above, hedonic evaluation mechanisms evolved to motivate adaptive behaviour and are conserved across species (e.g., *C. elegans*, *Homo sapiens*), scenarios and stimuli (e.g., food, sex, artworks). Active inference explains this continuum from first principles, with behaviour driven by the imperative to maintain preferred states shaped by evolutionary adaptations to an ecological niche. By integrating hedonic evaluation within active inference, hedonic foraging provides the computational mechanism through which reward maximisation drives and monitors behaviour in organisms with a reward system, situating the commonalities discussed above within a continuum of behaviours and species.

6. Implications and potential of the hedonic foraging framework

6.1. Contributions to the hedonic evaluation and active inference frameworks

We have presented hedonic foraging as a computational mechanism through which reward maximisation drives behaviour in organisms with a reward system, with hedonic evaluation monitoring and motivating free energy minimisation relative to specific, embodied affordances and conditions. The theoretical and mathematical integration of hedonic evaluation and active inference entails a reciprocal enrichment of both frameworks, advancing our understanding of how motivated behaviour aligns with reward-driven behaviour. In the remainder of this section, we succinctly discuss the implications of merging hedonic evaluation and active inference in hedonic foraging, and how this integrated approach addresses the central questions raised in the introduction.

Through hedonic foraging, active inference makes several key contributions to hedonic evaluation. Traditionally, hedonic evaluation has lacked a formal mechanism for quantifying and comparing hedonic values and translating these values into decisions and actions. In our proposal, active inference fills this gap by providing hedonic evaluation with a computational framework for understanding motivated behaviour. By formalising hedonic evaluation within the active inference framework, researchers can systematically examine how the interplay of factors such as internal states and environmental cues leads to behavioural outcomes. Crucially, hedonic foraging unifies *wanting* and *liking*, the core components of hedonic evaluation, as complementary facets of a single objective: maximising reward, that is, minimising VFE. Furthermore, this approach embraces the *common currency* hypothesis in that it acknowledges a single evaluative mechanism regardless of reward type, and transcends it by proposing a unified mechanism for behaviour across complexity and taxa grounded in free energy minimisation. In doing so, active inference provides hedonic evaluation with a first principle account of behaviour, linking neurobiological processes to computational goals.

Conversely, through our framework, hedonic evaluation offers significant insights with the potential to enhance and refine active inference. The mechanisms underlying evaluation and reward are increasingly well defined in active inference. Reward maximisation is integrated through prior preferences that encode desired outcomes. These preferences guide policy selection by minimising EFE, which balances exploration and exploitation. Additionally, active inference formalises reward-related behaviour by treating prediction errors and learning rates as precision-weighted processes, akin to mechanisms in reinforcement learning but grounded in a unifying Bayesian framework. Crucially, dopamine has emerged as a key candidate for encoding precision-weighted prediction errors in active inference. Phasic dopamine signals are thought to represent transient changes in precision, guiding learning and action by signalling reward prediction errors (Bogacz, 2020; Friston et al., 2012), while tonic dopamine signals (Bogacz, 2020), or the locus coeruleus–noradrenaline system (Sales et al., 2019), modulate the baseline confidence in policies, balancing exploratory and exploitative behaviour.

Despite these neurobiological formulations of active inference, much remains unknown about its precise neurobiological correlates. In particular, the neurobiological implementation of VFE minimisation is not so well formalised in active inference, and that of the correspondence between habits EFE and VFE has not yet been clarified. Thanks to its grounding in hedonic evaluation, the hedonic foraging framework outlines possible directions to address these gaps. As a first and paramount step, testing whether the activation of hedonic hotspots underlying

liking corresponds to a reduction of VFE in active inference is crucial to understanding how behavioural choices are evaluated. Let us consider, for example, an empirical task in which the organism has to choose among several options, such as a multi-armed bandit (Schwartenbeck et al., 2019). The choices leading to greater rewards correspond to a greater reduction in VFE. According to hedonic foraging, these choices would relate to the dopamine release signalling EFE and would subsequently be associated with the activation of hedonic hotspots. Testing such hypotheses would provide insights into the neurobiological correlates of active inference. Casting the constructs of hedonic evaluation—supported by its extensive neurobiological foundation (Berridge & Kringelbach, 2013, 2015)—within the active inference framework enables hedonic foraging to generate novel hypotheses and investigate the potential neurobiological mechanisms underlying free energy minimisation.

In summary, hedonic foraging fosters a mutual enrichment of hedonic evaluation—the fundamental adaptive neurobiological process of assessing stimuli as beneficial or harmful—and active inference—a unified, principled theory of behaviour. In doing so, hedonic foraging has the potential to explain how information encoded in hedonic values is integrated and transformed into motivated behaviour, unifying computational and neurobiological mechanisms of behaviour into a single theoretical framework.

6.2. Explaining and predicting empirical data

The value of any theory or model resides in its heuristic utility and predictive power—its ability to explain existing phenomena and anticipate new data (Lipton, 2005). The central concept for understanding how hedonic foraging can address empirical questions is *computational phenotyping*. At the core of hedonic foraging, regardless of the scenario under scrutiny, lies a generative model, which establishes a mapping between model parameters and behaviour. Computational phenotyping is the process of retrieving model parameters from empirical data by inverting such mapping (Schwartenbeck & Friston, 2016). Computational phenotyping has emerged as a powerful tool for characterising individual variability in domains such as social behaviour (Xiang et al., 2012), personality, developmental and clinical neuroscience (Patzelt et al., 2018), cognition (Schurr et al., 2024) and genetics (Drouin et al., 2016). It has also been applied to active inference, particularly in computational psychiatry (Schwartenbeck & Friston, 2016) and clinical neuroscience (Smith et al., 2021), offering a dual benefit: First, it harnesses the power of active inference—with its principled formulation and parameters interpretable in terms of cognitive processes—to provide a deeper understanding of cognitive phenomena. Second, it enables the simulation and prediction of behaviour based on empirically derived parameters.

Hedonic foraging promotes the application of computational phenotyping to empirical experiments involving hedonic evaluation. To illustrate the potential benefits, let us consider an example from behavioural and cognitive ecology, where measuring preference and choice is central (Rosenthal, 2017). Predators preferentially choose cryptically coloured over conspicuous prey, as the latter are often associated with defences such as toxins that render them unpalatable (Mappes et al., 2005; Penacchio et al., 2024). This preference aligns with the imperative to minimise surprise in active inference, as sampling conspicuous prey can be associated with surprise (Penacchio et al., 2025). However, the preference for cryptic prey can be reduced or reversed when food is scarce, as the imperative to avoid hunger takes precedence (Barnett et al., 2007). Analysing behaviour in the context of predation—and other domains such as sexual selection—through the lens of hedonic foraging leverages deep

neurophysiological knowledge from hedonic evaluation within a principled computational framework, and provides insights into individual differences, including personality traits, sex and species-specific behaviours.

Another relevant instance is the study of music appreciation, a model domain for a broad investigation of cognition and behaviour. Listening to music involves the interplay of factors such as individual brain structure (Loui et al., 2017) and function (Martínez-Molina et al., 2016), physiological state (Dreher et al., 2007; Thomas et al., 2015), individual susceptibility in general (Aydogan et al., 2018) and specific to framing effects (Lopez-Persem et al., 2016), personality (Rentfrow & Gosling, 2003), expertise (Alluri et al., 2015; Brattico et al., 2016), preferences (Palmer & Griscom, 2013; Smith & Melara, 1990), familiarity with the piece (Pereira et al., 2011; van den Bosch et al., 2013), dynamic and schematic (stylistic) expectations (Huron, 2006; Pearce, 2018) about the predictability and uncertainty of the musical structures (Cheung et al., 2019; Gold et al., 2019), physical activity (North & Hargreaves, 2000), needs and goals (Sachs et al., 2015; Thoma et al., 2012) and social setting (Curzel et al., 2024; Nummenmaa et al., 2021). The reward system integrates these inputs and generates value signals through dopamine release and activation of hedonic hotspots (Zatorre, 2015; Zatorre & Salimpoor, 2013). This activity triggers action programs (Damasio & Carvalho, 2013) such as those associated with musical chills—elevating heart rate, skin conductance and respiration (de Fleurian & Pearce, 2021)—and policies directed to continuing or ceasing the listening experience. The hedonic impact of the experience contributes to the formation of preferences (Finnas, 1989; Kathios et al., 2024) and habits, thus informing subsequent behaviour—e.g., listening to the piece again, seeking related works, exploring new styles, learning to play an instrument, composing a piece and even devoting professionally to music.

By applying computational phenotyping to these scenarios, researchers can disentangle how the complex interplay between habit, wanting and liking shapes behaviour. This approach allows experiments to move beyond descriptive measures, uncovering deeper mechanistic explanations that integrate the multiplicity of factors driving behaviour.

Some of these factors correspond to general elements of the generative models, such as learning rates or the balance between exploration and exploitation, which apply to multiple scenarios. Curiosity (Schutte & Malouff, 2020), particularly its epistemic form (Hardy et al., 2017), can be mapped onto lower precision (less determinacy) in policy selection (β in **Box 2,** (Schwartenbeck et al., 2019). As shown in **Scenario 3**, curiosity may counteract and even reverse habits and biases. In **Scenario 4**, we demonstrated how the interplay between abilities and preferences drives behaviour and how these can be learned through reward maximisation.

Other factors are more specific, enabling testing structured hypotheses. This is the case of individual traits that entail specific mappings in the generative models. Expertise (Ullén et al., 2016) constitutes a prominent exemplary trait to test specific hypotheses through computational phenotyping, as it can be characterised through a more deterministic (less stochastic) mapping between states and observations ($\boldsymbol{A}$ in **Box 2**) and more deterministic transitions between states over time (e.g., what the resolution of a harmonic sequence can be) ($\boldsymbol{B}$ in **Box 2**), as a result of extensive statistical learning. For instance, individual sound events, their voice leading and tonal relationships in a Western classical musical piece are more precisely defined for classical musicians than aficionados (Schlaug, 2015; Zatorre et al., 2007). In active inference, this precision can be expressed as a higher probability of correspondence between the actual sounds and their perception in the generative model's tensor $\boldsymbol{A}$. Similarly, the probabilities for predicting the next event within a specific set in tensor $\boldsymbol{B}$ are significantly higher for trained musicians, reflecting their enhanced ability to anticipate musical structure. This explains why

musicians experience greater surprise when tonal or voice-leading relationships are violated; that is, why they are less tolerant to mistakes (Hernández et al., 2019).

In summary, computational phenotyping within the hedonic foraging framework offers a principled and versatile approach to studying motivated behaviour. The parameters inferred from observed behaviour can be meaningfully interpreted in terms of fundamental cognitive processes, drawing on core principles from hedonic evaluation as well as foundational constructs within active inference. Moreover, the theory facilitates the generation of testable predictions by simulating behaviour based on specific parameter values, which can then be compared to empirical observations in experimental settings.

7. Conclusion

In this paper, we propose a natural operationalisation of hedonic values in behaviour by mapping the phases of hedonic evaluation and their corresponding value signals onto the core components of active inference. In our interpretation, VFE corresponds to hedonic impact or core liking, quantifying sensory pleasure—or displeasure—of ongoing experiences, while EFE captures the motivational drive or core wanting to pursue rewarding outcomes. This synthesis forms the foundation of our theoretical framework, hedonic foraging, which describes how hedonic evaluation generates and evaluates motivated behaviour.

Grounded in active inference, hedonic foraging provides a unifying framework for understanding preference-seeking behaviours through principles akin to those governing physical systems. Drawing on analogies from statistical physics, it conceptualises organisms as agents navigating probabilistic spaces to minimise VFE—similar to particles minimising a functional such as the Lagrangian (Friston, 2013; Friston et al., 2023). In this view, hedonic evaluation guides organisms along behavioural trajectories that minimise surprise, mirroring the principle of least action in physics. Crucially, this predictive framework relies on an embodied model of the world—shaped by an organism s ecological niche—that generates expectations about and evaluates the outcomes of actions in the form of hedonic values.

A key strength of hedonic foraging is its ability to unify diverse preference-seeking behaviours under a shared mechanistic framework. This extends the common currency hypothesis beyond static neural correlates of hedonic value to highlight the dynamic mechanisms underpinning hedonic choice in action planning. In doing so, our proposal bridges fundamental biological processes with higher-order cultural endeavours typically considered uniquely human, suggesting shared roots in adaptive mechanisms.

Ultimately, hedonic foraging reinforces the idea that cognition is not primarily about abstract reasoning or knowledge accumulation but about engaging in adaptive interactions with the environment and its affordances. This perspective aligns with the growing recognition that purposive interaction is key to understanding cognition (Pezzulo et al., 2024) and represents a promising avenue for advancing artificial intelligence. The move from traditional milestones in the field, such as the Turing test to the embodied Turing test (Zador et al., 2023), reflects this paradigm shift. We hope our proposal will inspire further investigation into the neurobiology of motivated behaviour, bridging foundational biological mechanisms with complex cognitive functions and fostering interdisciplinary scientific progress.

Bibliography

Alluri, V., Brattico, E., Toiviainen, P., Burunat, I., Bogert, B., Numminen, J., & Kliuchko, M. (2015). Musical expertise modulates functional connectivity of limbic regions during continuous music listening. *Psychomusicology: Music, Mind, and Brain*, *25*(4), 443-454. https://doi.org/10.1037/pmu0000124

Anderson, D. J., & Adolphs, R. (2014). A Framework for Studying Emotions across Species [Review]. *Cell*, *157*(1), 187-200. https://doi.org/10.1016/j.cell.2014.03.003

Aydogan, G., Flaig, N., Ravi, S. N., Large, E. W., McClure, S. M., & Margulis, E. H. (2018). Overcoming Bias: Cognitive Control Reduces Susceptibility to Framing Effects in Evaluating Musical Performance [Article]. *Scientific Reports*, *8*, 9, Article 6229. https://doi.org/10.1038/s41598-018-24528-3

Barnett, C. A., Bateson, M., & Rowe, C. (2007). State-dependent decision making: educated predators strategically trade off the costs and benefits of consuming aposematic prey [Article]. *Behavioral Ecology*, *18*(4), 645-651. https://doi.org/10.1093/beheco/arm027

Barrett, L. F. (2017). The theory of constructed emotion: an active inference account of interoception and categorization [Article]. *Social Cognitive and Affective Neuroscience*, *12*(1), 1-23. https://doi.org/10.1093/scan/nsw154

Barron, A. B., Sovik, E., & Cornish, J. L. (2010). The roles of dopamine and related compounds in reward-seeking behavior across animal phyla [Review]. *Frontiers in Behavioral Neuroscience*, *4*, 9, Article 163. https://doi.org/10.3389/fnbeh.2010.00163

Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value [Article]. *Neuroimage*, *76*(1), 412-427. https://doi.org/10.1016/j.neuroimage.2013.02.063

Berridge, K. C. (1991). Modulation of taste affect by hunger, caloric satiety, and sensory-specific satiety in the rat [Article]. *Appetite*, *16*(2), 103-120. https://doi.org/10.1016/0195-6663(91)90036-r

Berridge, K. C. (2004). Motivation concepts in behavioral neuroscience [Review]. *Physiology and Behavior*, *81*(2), 179-209. https://doi.org/10.1016/j.physbeh.2004.02.004

Berridge, K. C., & Dayan, P. (2021). Liking [Editorial Material]. *Current Biology*, *31*(24), R1555-R1555. https://doi.org/10.1016/j.cub.2021.09.069

Berridge, K. C., & Kringelbach, M. L. (2008). Affective neuroscience of pleasure: reward in humans and animals [Review]. *Psychopharmacology*, *199*(3), 457-480. https://doi.org/10.1007/s00213-008-1099-6

Berridge, K. C., & Kringelbach, M. L. (2013). Neuroscience of affect: brain mechanisms of pleasure and displeasure [Review]. *Current Opinion in Neurobiology*, *23*(3), 294-303. https://doi.org/10.1016/j.conb.2013.01.017

Berridge, K. C., & Kringelbach, M. L. (2015). Pleasure Systems in the Brain [Review]. *Neuron*, *86*(3), 646-664. https://doi.org/10.1016/j.neuron.2015.02.018

Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? [Review]. *Brain Research Reviews*, *28*(3), 309-369. https://doi.org/10.1016/s0165-0173(98)00019-8

Birch, L. L., & Marlin, D. W. (1982). I don't like it; I never tried it: Effects of exposure on two-year-old children's food preferences. *Appetite*, *3*(4), 353-360. https://doi.org/https://doi.org/10.1016/S0195-6663(82)80053-6

Bogacz, R. (2020). Dopamine role in learning and action inference [Article]. *Elife*, *9*, 33, Article e53262. https://doi.org/10.7554/eLife.53262

Bornstein, R. F., & Dagostino, P. R. (1992). Stimulus-recognition and the mere exposure effect [Article]. *Journal of Personality and Social Psychology*, *63*(4), 545-552. https://doi.org/10.1037//0022-3514.63.4.545

Brattico, E., Bogert, B., Alluri, V., Tervaniemi, M., Eerola, T., & Jacobsen, T. (2016). It's Sad but I Like It: The Neural Dissociation Between Musical Emotions and Liking in Experts and Laypersons [Article]. *Frontiers in Human Neuroscience*, *9*, 21, Article 676. https://doi.org/10.3389/fnhum.2015.00676

Brieber, D., Nadal, M., & Leder, H. (2015). In the white cube: Museum context enhances the valuation and memory of art [Article]. *Acta Psychologica*, *154*, 36-42. https://doi.org/10.1016/j.actpsy.2014.11.004

Brielmann, A. A., & Dayan, P. (2022). A Computational Model of Aesthetic Value [Article]. *Psychological Review*, *129*(6), 1319-1337. https://doi.org/10.1037/rev0000337

Calhoun, A. J., Chalasani, S. H., & Sharpee, T. O. (2014). Maximally informative foraging by Caenorhabditis elegans [Article]. *Elife*, *3*, 13, Article e04220. https://doi.org/10.7554/eLife.04220

Cannon, W. B. (1929). Organization for physiological homeostasis [Review]. *Physiological Reviews*, *9*(3), 399-431. https://doi.org/10.1152/physrev.1929.9.3.399

Cheung, V. K. M., Harrison, P. M. C., Meyer, L., Pearce, M. T., Haynes, J. D., & Koelsch, S. (2019). Uncertainty and Surprise Jointly Predict Musical Pleasure and Amygdala, Hippocampus, and Auditory Cortex Activity [Article]. *Current Biology*, *29*(23), 4084-+. https://doi.org/10.1016/j.cub.2019.09.067

Chiel, H. J., & Beer, R. D. (1997). The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment [Editorial Material]. *Trends in Neurosciences*, *20*(12), 553-557. https://doi.org/10.1016/s0166-2236(97)01149-1

Chudasama, Y. (2011). Animal Models of Prefrontal-Executive Function [Article]. *Behavioral Neuroscience*, *125*(3), 327-343. https://doi.org/10.1037/a0023766

Cisek, P. (2022). Evolution of behavioural control from chordates to primates. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *377*(1844), 20200522. https://doi.org/doi:10.1098/rstb.2020.0522

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science [Article]. *Behavioral and Brain Sciences*, *36*(3), 181-204. https://doi.org/10.1017/s0140525x12000477

Conant, R. C., & Ross Ashby, W. (1970). Every good regulator of a system must be a model of that system †. *International Journal of Systems Science*, *1*(2), 89-97. https://doi.org/10.1080/00207727008920220

Curzel, F., Osiurak, F., Trân, E., Tillmann, B., Ripollés, P., & Ferreri, L. (2024). Enhancing musical pleasure through shared musical experience [Article]. *Iscience*, *27*(6), 17, Article 109964. https://doi.org/10.1016/j.isci.2024.109964

Damasio, A., & Carvalho, G. B. (2013). OPINION The nature of feelings: evolutionary and neurobiological origins [Review]. *Nature Reviews Neuroscience*, *14*(2), 143-152. https://doi.org/10.1038/nrn3403

Darwin, C. (1872). *The expression of the emotions in man and animals* [doi:10.1037/10001-000]. John Murray. https://doi.org/10.1037/10001-000

Dayan, P. (2009). Goal-directed control and its antipodes [Article; Proceedings Paper]. *Neural Networks*, *22*(3), 213-219. https://doi.org/10.1016/j.neunet.2009.03.004

Dayan, P. (2022). "Liking" as an early and editable draft of long-run affective value [Review]. *PLoS Biology*, *20*(1), 15, Article e3001476. https://doi.org/10.1371/journal.pbio.3001476

de Fleurian, R., & Pearce, M. T. (2021). Chills in Music: A Systematic Review [Review]. *Psychological Bulletin*, *147*(9), 890-920. https://doi.org/10.1037/bul0000341

Dolan, R. J., & Dayan, P. (2013). Goals and Habits in the Brain [Review]. *Neuron*, *80*(2), 312-325. https://doi.org/10.1016/j.neuron.2013.09.007

Doya, K., Ishii, S., Pouget, A., & Rao, R. P. N. (2007). *Bayesian brain: Probabilistic approaches to neural coding* (K. Doya, S. Ishii, A. Pouget, & R. P. N. Rao, Eds.). MIT Press.

Dreher, J. C., Schmidt, P. J., Kohn, P., Furman, D., Rubinow, D., & Berman, K. F. (2007). Menstrual cycle phase modulates reward-related neural function in women [Article]. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(7), 2465-2470. https://doi.org/10.1073/pnas.0605569104

Drouin, A., Giguère, S., Déraspe, M., Marchand, M., Tyers, M., Loo, V. G., Bourgault, A. M., Laviolette, F., & Corbeil, J. (2016). Predictive computational phenotyping and biomarker discovery using reference-free genome comparisons [Article]. *BMC Genomics*, *17*, 15, Article 754. https://doi.org/10.1186/s12864-016-2889-6

Duckworth, R. A. (2009). The role of behavior in evolution: a search for mechanism [Review]. *Evolutionary Ecology*, *23*(4), 513-531. https://doi.org/10.1007/s10682-008-9252-6

Finnas, L. (1989). HOW CAN MUSICAL PREFERENCES BE MODIFIED - A RESEARCH REVIEW [Review]. *Bulletin of the Council for Research in Music Education*(102), 1-58. <Go to ISI>://WOS:A1989AR30900001

FitzGerald, T. H. B., Dolan, R. J., & Friston, K. (2015). Dopamine, reward learning, and active inference [Article]. *Frontiers in Computational Neuroscience*, *9*, 16, Article 136. https://doi.org/10.3389/fncom.2015.00136

Friedman, N. P., & Robbins, T. W. (2022). The role of prefrontal cortex in cognitive control and executive function [Review]. *Neuropsychopharmacology*, *47*(1), 72-89. https://doi.org/10.1038/s41386-021-01132-0

Friston, K. (2013). Life as we know it [Article]. *Journal of the Royal Society Interface*, *10*(86), 12, Article 20130475. https://doi.org/10.1098/rsif.2013.0475

Friston, K., Da Costa, L., Sakthivadivel, D. A. R., Heins, C., Pavliotis, G. A., Ramstead, M., & Parr, T. (2023). Path integrals, particular kinds, and strange things [Review]. *Physics of Life Reviews*, *47*, 35-62. https://doi.org/10.1016/j.plrev.2023.08.016

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O'Doherty, J., & Pezzulo, G. (2016). Active inference and learning [Review]. *Neuroscience and Biobehavioral Reviews*, *68*, 862-879. https://doi.org/10.1016/j.neubiorev.2016.06.022

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active Inference: A Process Theory [Article]. *Neural Computation*, *29*(1), 1-49. https://doi.org/10.1162/NECO_a_00912

Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value [Article]. *Cognitive Neuroscience*, *6*(4), 187-214. https://doi.org/10.1080/17588928.2015.1020053

Friston, K. J. (2010). The free-energy principle: a unified brain theory? [Review]. *Nature Reviews Neuroscience*, *11*(2), 127-138. https://doi.org/10.1038/nrn2787

Friston, K. J., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain [Article]. *Journal of Physiology-Paris*, *100*(1-3), 70-87. https://doi.org/10.1016/j.jphysparis.2006.10.001

Friston, K. J., Lin, M., Frith, C. D., Pezzulo, G., Hobson, J. A., & Ondobaka, S. (2017). Active Inference, Curiosity and Insight [Article]. *Neural Computation*, *29*(10), 2633-2683. https://doi.org/10.1162/NECO_a_00999

Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., Dolan, R. J., Moran, R., Stephan, K. E., & Bestmann, S. (2012). Dopamine, Affordance and Active Inference [Article]. *PLoS Computational Biology*, *8*(1), 20, Article e1002327. https://doi.org/10.1371/journal.pcbi.1002327

Gallagher, S. (2017). *Enactivist Interventions: Rethinking the Mind*. Oxford University Press. https://doi.org/10.1093/oso/9780198794325.001.0001

Gartus, A., & Leder, H. (2014). The white cube of the museum versus the gray cube of the street: The role of context in aesthetic evaluations. *Psychology of Aesthetics, Creativity, and the Arts*, *8*(3), 311-320. https://doi.org/10.1037/a0036847
Gibson, J. J. (1986). *The Ecological Approach to Visual Perception*. Lawrence Erlbaum Associates. https://books.google.fr/books?id=DrhCCWmJpWUC
Gintis, H. (2007). A framework for the unification of the behavioral sciences [Review]. *Behavioral and Brain Sciences*, *30*(1), 1-+. https://doi.org/10.1017/s0140525x07000581
Gold, B. P., Pearce, M. T., Mas-Herrero, E., Dagher, A., & Zatorre, R. J. (2019). Predictability and Uncertainty in the Pleasure of Music: A Reward for Learning? [Article]. *Journal of Neuroscience*, *39*(47), 9397-9409. https://doi.org/10.1523/jneurosci.0428-19.2019
Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making [Review; Book Chapter]. *Annual Review of Neuroscience*, *30*, 535-574. https://doi.org/10.1146/annurev.neuro.29.051605.113038
Gottlieb, J., & Oudeyer, P. Y. (2018). Towards a neuroscience of active sampling and curiosity [Review]. *Nature Reviews Neuroscience*, *19*(12), 758-770. https://doi.org/10.1038/s41583-018-0078-0
Gregory, R. L. (1980). Perceptions as hypotheses [Article]. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, *290*(1038), 181-197. https://doi.org/10.1098/rstb.1980.0090
Güntürkün, O. (2005). The avian 'prefrontal cortex' and cognition [Review]. *Current Opinion in Neurobiology*, *15*(6), 686-693. https://doi.org/10.1016/j.conb.2005.10.003
Güntürkün, O., & Bugnyar, T. (2016). Cognition without Cortex [Review]. *Trends in Cognitive Sciences*, *20*(4), 291-303. https://doi.org/10.1016/j.tics.2016.02.001
Hardy, J. H., Ness, A. M., & Mecca, J. (2017). Outside the box: Epistemic curiosity as a predictor of creative problem solving and creative performance [Article]. *Personality and Individual Differences*, *104*, 230-237. https://doi.org/10.1016/j.paid.2016.08.004
Hernández, M., Palomar-García, M. A., Nohales-Nieto, B., Olcina-Sempere, G., Villar-Rodríguez, E., Pastor, R., Avila, C., & Parcet, M. A. (2019). Separate Contribution of Striatum Volume and Pitch Discrimination to Individual Differences in Music Reward [Article]. *Psychological Science*, *30*(9), 1352-1361, Article 0956797619859339. https://doi.org/10.1177/0956797619859339
Hesp, C., Smith, R., Parr, T., Allen, M., Friston, K. J., & Ramstead, M. J. D. (2021). Deeply Felt Affect: The Emergence of Valence in Deep Active Inference [Article]. *Neural Computation*, *33*(2), 398-446. https://doi.org/10.1162/neco_a_01341
Huron, D. (2006). *Sweet anticipation: Music and the psychology of expectation*. The MIT Press.
Jarvis, E. D., Gunturkun, O., Bruce, L., Csillag, A., Karten, H., Kuenzel, W., Medina, L., Paxinos, G., Perkel, D. J., Shimizu, T., Striedter, G., Wild, J. M., Ball, G. F., Dugas-Ford, J., Durand, S. E., Hough, G. E., Husband, S., Kubikova, L., Lee, D. W.,…Avian Brain Nomenclature, C. (2005). Avian brains and a new understanding of vertebrate brain evolution [Review]. *Nature Reviews Neuroscience*, *6*(2), 151-159. https://doi.org/10.1038/nrn1606
Joffily, M., & Coricelli, G. (2013). Emotional Valence and the Free-Energy Principle [Article]. *PLoS Computational Biology*, *9*(6), 14, Article e1003094. https://doi.org/10.1371/journal.pcbi.1003094
Kathios, N., Sachs, M. E., Zhang, E., Ou, Y. T., & Loui, P. (2024). Generating New Musical Preferences From Multilevel Mapping of Predictions to Reward [Article]. *Psychological Science*, *35*(1), 34-54. https://doi.org/10.1177/09567976231214185
Kidd, C., & Hayden, B. Y. (2015). The Psychology and Neuroscience of Curiosity [Review]. *Neuron*, *88*(3), 449-460. https://doi.org/10.1016/j.neuron.2015.09.010
Kleckner, I. R., Zhang, J. H., Touroutoglou, A., Chanes, L., Xia, C. J., Simmons, W. K., Quigley, K. S., Dickerson, B. C., & Barrett, L. F. (2017). Evidence for a large-scale brain system

supporting allostasis and interoception in humans [Article]. *Nature Human Behaviour*, *1*(5), 14, Article 0069. https://doi.org/10.1038/s41562-017-0069
Klink, P. C., van Wezel, R. J. A., & van Ee, R. (2012). United we sense, divided we fail: context-driven perception of ambiguous visual stimuli [Review]. *Philosophical Transactions of the Royal Society B-Biological Sciences*, *367*(1591), 932-941. https://doi.org/10.1098/rstb.2011.0358
Koelsch, S., Bashevkin, T., Kristensen, J., Tvedt, J., & Jentschke, S. (2019). Heroic music stimulates empowering thoughts during mind-wandering [Article]. *Scientific Reports*, *9*, 10, Article 10317. https://doi.org/10.1038/s41598-019-46266-w
Koelsch, S., Vuust, P., & Friston, K. (2019). Predictive Processes and the Peculiar Case of Music [Review]. *Trends in Cognitive Sciences*, *23*(1), 63-77. https://doi.org/10.1016/j.tics.2018.10.006
Kording, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control [Article; Proceedings Paper]. *Trends in Cognitive Sciences*, *10*(7), 319-326. https://doi.org/10.1016/j.tics.2006.05.003
Kounios, J., & Beeman, M. (2014). The Cognitive Neuroscience of Insight. In S. T. Fiske (Ed.), *Annual Review of Psychology, Vol 65* (Vol. 65, pp. 71-93). Annual Reviews. https://doi.org/10.1146/annurev-psych-010213-115154
Kringelbach, M. L. (2005). The human orbitofrontal cortex: Linking reward to hedonic experience [Review]. *Nature Reviews Neuroscience*, *6*(9), 691-702. https://doi.org/10.1038/nrn1747
Leder, H. (2003). Familiar and Fluent! Style-Related Processing Hypotheses in Aesthetic Appreciation. *Empirical Studies of the Arts*, *21*(2), 165-175. https://doi.org/10.2190/g6mk-6kl7-ketb-elnd
LeDoux, J. E. (2012). Evolution of human emotion: A view through fear. In M. A. Hofman & D. Falk (Eds.), *Evolution of the Primate Brain: From Neuron to Behavior* (Vol. 195, pp. 431-442). Elsevier Science Bv. https://doi.org/10.1016/b978-0-444-53860-4.00021-0
Levy, D. J., & Glimcher, P. W. (2012). The root of all value: a neural common currency for choice [Review]. *Current Opinion in Neurobiology*, *22*(6), 1027-1038. https://doi.org/10.1016/j.conb.2012.06.001
Lim, J. Y. (2011). Hedonic scaling: A review of methods and theory [Review]. *Food Quality and Preference*, *22*(8), 733-747. https://doi.org/10.1016/j.foodqual.2011.05.008
Lipton, P. (2005). Testing hypotheses: Prediction and prejudice [Review]. *Science*, *307*(5707), 219-221. https://doi.org/10.1126/science.1103024
Lopez, J. C., & Lopez, D. (1985). Killer whales (Orcinus orca) of Patagonia, and their behavior of intentional stranding while hunting nearshore [Note]. *Journal of Mammalogy*, *66*(1), 181-183. https://doi.org/10.2307/1380981
López-Gambero, A. J., Martínez, F., Salazar, K., Cifuentes, M., & Nualart, F. (2019). Brain Glucose-Sensing Mechanism and Energy Homeostasis [Review]. *Molecular Neurobiology*, *56*(2), 769-796. https://doi.org/10.1007/s12035-018-1099-4
Lopez-Persem, A., Domenech, P., & Pessiglione, M. (2016). How prior preferences determine decision-making frames and biases in the human brain [Article]. *Elife*, *5*, 20, Article e20317. https://doi.org/10.7554/eLife.20317
Loui, P., Patterson, S., Sachs, M. E., Leung, Y., Zeng, T. M., & Przysinda, E. (2017). White Matter Correlates of Musical Anhedonia: Implications for Evolution of Music. *Frontiers in Psychology*, *8*, Article 1664. https://doi.org/10.3389/fpsyg.2017.01664
Madirolas, G., Al-Asmar, A., Gaouar, L., Marie-Louise, L., Garza-Enríquez, A., Rodríguez-Rada, V., Khona, M., Dal Bello, M., Ratzke, C., Gore, J., & Pérez-Escudero, A. (2023). Caenorhabditis elegans foraging patterns follow a simple rule of thumb [Article]. *Communications Biology*, *6*(1), 13, Article 841. https://doi.org/10.1038/s42003-023-05220-3

Malekzadeh, P., & Plataniotis, K. N. (2024). Active Inference and Reinforcement Learning: A Unified Inference on Continuous State and Action Spaces Under Partial Observability [Article]. *Neural Computation*, *36*(10), 2073-2135. https://doi.org/10.1162/neco_a_01698

Mappes, J., Marples, N., & Endler, J. A. (2005). The complex business of survival by aposematism [Review]. *Trends in Ecology & Evolution*, *20*(11), 598-603. https://doi.org/10.1016/j.tree.2005.07.011

Martínez-Molina, N., Mas-Herrero, E., Rodríguez-Fornells, A., Zatorre, R. J., & Marco-Pallarés, J. (2016). Neural correlates of specific musical anhedonia [Article]. *Proceedings of the National Academy of Sciences of the United States of America*, *113*(46), E7337-E7345. https://doi.org/10.1073/pnas.1611211113

Marty, N., Dallaporta, M., & Thorens, B. (2007). Brain glucose sensing, counterregulation, and energy homeostasis [Review]. *Physiology*, *22*, 241-251. https://doi.org/10.1152/physiol.00010.2007

Mesoudi, A., Whiten, A., & Laland, K. N. (2006). Towards a unified science of cultural evolution [Review]. *Behavioral and Brain Sciences*, *29*(4), 329-+. https://doi.org/10.1017/s0140525x06009083

Montague, P. R., & King-Casas, B. (2007). Efficient statistics, common currencies and the problem of reward-harvesting [Review]. *Trends in Cognitive Sciences*, *11*(12), 514-519. https://doi.org/10.1016/j.tics.2007.10.002

Nadal, M., & Skov, M. (2024). The sensory valuation account of aesthetic experience [Review; Early Access]. *Nature Reviews Psychology*, 15. https://doi.org/10.1038/s44159-024-00385-y

Nguyen, D., Naffziger, E. E., & Berridge, K. C. (2021). Positive affect: nature and brain bases of liking and wanting [Article]. *Current Opinion in Behavioral Sciences*, *39*, 72-78. https://doi.org/10.1016/j.cobeha.2021.02.013

North, A. C., & Hargreaves, D. J. (2000). Musical preferences during and after relaxation and exercise [Article]. *American Journal of Psychology*, *113*(1), 43-67. https://doi.org/10.2307/1423460

Nummenmaa, L., Putkinen, V., & Sams, M. (2021). Social pleasures of music [Article]. *Current Opinion in Behavioral Sciences*, *39*, 196-202. https://doi.org/10.1016/j.cobeha.2021.03.026

O'Connell, L. A., & Hofmann, H. A. (2011). The Vertebrate mesolimbic reward system and social behavior network: A comparative synthesis [Review]. *Journal of Comparative Neurology*, *519*(18), 3599-3639. https://doi.org/10.1002/cne.22735

Palmer, S. E., & Griscom, W. S. (2013). Accounting for taste: Individual differences in preference for harmony [Article]. *Psychonomic Bulletin & Review*, *20*(3), 453-461. https://doi.org/10.3758/s13423-012-0355-2

Panksepp, J. (1998). The periconscious substrates of consciousness: Affective states and the evolutionary origins of the self. *Journal of Consciousness Studies*, *5*(5-6), 566-582.

Parr, T., Pezzulo, G., & Friston, K. J. (2022). *Active inference: the free energy principle in mind, brain, and behavior*. MIT Press.

Patzelt, E. H., Hartley, C. A., & Gershman, S. J. (2018). Computational Phenotyping: Using Models to Understand Individual Differences in Personality, Development, and Mental Illness. *Personal Neurosci*, *1*, e18. https://doi.org/10.1017/pen.2018.14

Pearce, M. T. (2018). Statistical learning and probabilistic prediction in music cognition: mechanisms of stylistic enculturation [Article; Proceedings Paper]. *Annals of the New York Academy of Sciences*, *1423*(1), 378-395. https://doi.org/10.1111/nyas.13654

Penacchio, O., Halpin, C. G., Cuthill, I. C., Lovell, P. G., Wheelwright, M., Skelhorn, J., Rowe, C., & Harris, J. M. (2024). A computational neuroscience framework for quantifying warning signals [Article]. *Methods in Ecology and Evolution*, *15*(1), 103-116. https://doi.org/10.1111/2041-210x.14268

Penacchio, O., Hämäläinen, L., Rojas, B., Summers, K., Yeager, J., Sherratt, T. N., & Exnerová, A. (2025). Cognitive ecology of surprise in predator-prey interactions [Review; Early Access]. *Functional Ecology*, 17. https://doi.org/10.1111/1365-2435.14750

Pereira, C. S., Teixeira, J., Figueiredo, P., Xavier, J., Castro, S. L., & Brattico, E. (2011). Music and Emotions in the Brain: Familiarity Matters [Article]. *PloS One*, *6*(11), 9, Article e27241. https://doi.org/10.1371/journal.pone.0027241

Perry, C. J., & Barron, A. B. (2013). Neural Mechanisms of Reward in Insects. In M. R. Berenbaum (Ed.), *Annual Review of Entomology, Vol 58* (Vol. 58, pp. 543-562). Annual Reviews. https://doi.org/10.1146/annurev-ento-120811-153631

Pessiglione, M., & Lebreton, M. (2015). From the Reward Circuit to the Valuation System: How the Brain Motivates Behavior. In G. H. E. Gendolla, M. Tops, & S. L. Koole (Eds.), *Handbook of Biobehavioral Approaches to Self-Regulation* (pp. 157-173). Springer New York. https://doi.org/10.1007/978-1-4939-1236-0_11

Pezzulo, G., & Cisek, P. (2016). Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition [Review]. *Trends in Cognitive Sciences*, *20*(6), 414-424. https://doi.org/10.1016/j.tics.2016.03.013

Pezzulo, G., Parr, T., & Friston, K. (2024). Active inference as a theory of sentient behavior [Article]. *Biological Psychology*, *186*, 9, Article 108741. https://doi.org/10.1016/j.biopsycho.2023.108741

Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control [Review]. *Progress in Neurobiology*, *134*, 17-35. https://doi.org/10.1016/j.pneurobio.2015.09.001

Pitman, R. L., & Durban, J. W. (2012). Cooperative hunting behavior, prey selectivity and prey handling by pack ice killer whales (Orcinus orca), type B, in Antarctic Peninsula waters [Article]. *Marine Mammal Science*, *28*(1), 16-36. https://doi.org/10.1111/j.1748-7692.2010.00453.x

Pliner, P. (1982). The effects of mere exposure on liking for edible substances [Article]. *Appetite*, *3*(3), 283-290. https://doi.org/10.1016/s0195-6663(82)80026-3

Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: knowns and unknowns [Review]. *Nature Neuroscience*, *16*(9), 1170-1178. https://doi.org/10.1038/nn.3495

Ramsay, D. S., & Woods, S. C. (2014). Clarifying the Roles of Homeostasis and Allostasis in Physiological Regulation [Article]. *Psychological Review*, *121*(2), 225-247. https://doi.org/10.1037/a0035942

Ramstead, M. J. D., Kirchhoff, M. D., & Friston, K. J. (2020). A tale of two densities: active inference is enactive inference [Article]. *Adaptive Behavior*, *28*(4), 225-239, Article 1059712319862774. https://doi.org/10.1177/1059712319862774

Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making [Review]. *Nature Reviews Neuroscience*, *9*(7), 545-556. https://doi.org/10.1038/nrn2357

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects [Article]. *Nature Neuroscience*, *2*(1), 79-87. https://doi.org/10.1038/4580

Reber, R., Schwarz, N., & Winkielman, P. (2004). Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? [Review]. *Personality and Social Psychology Review*, *8*(4), 364-382. https://doi.org/10.1207/s15327957pspr0804_3

Reber, R., Winkielman, P., & Schwarz, N. (1998). Effects of perceptual fluency on affective judgments [Article]. *Psychological Science*, *9*(1), 45-48. https://doi.org/10.1111/1467-9280.00008

Rentfrow, P. J., & Gosling, S. D. (2003). The do re Mi's of everyday life: The structure and personality correlates of music preferences [Article]. *Journal of Personality and Social Psychology*, *84*(6), 1236-1256. https://doi.org/10.1037/0022-3514.84.6.1236
Roh, E., Song, D. K., & Kim, M. S. (2016). Emerging role of the brain in the homeostatic regulation of energy and glucose metabolism [Review]. *Experimental and Molecular Medicine*, *48*, 12, Article e216. https://doi.org/10.1038/emm.2016.4
Rosenthal, G. G. (2017). *Mate choice: The evolution of sexual decision making from microbes to humans*. Princeton University Press.
Ross, S., & Hansen, N. C. (2016). Dissociating Prediction Failure: Considerations from Music Perception [Editorial Material]. *Journal of Neuroscience*, *36*(11), 3103-3105. https://doi.org/10.1523/jneurosci.0053-16.2016
Roth, W. M., & Jornet, A. (2013). Situated cognition [Article]. *Wiley Interdisciplinary Reviews-Cognitive Science*, *4*(5), 463-478. https://doi.org/10.1002/wcs.1242
Sachs, M. E., Damasio, A., & Habibi, A. (2015). The pleasures of sad music: a systematic review [Review]. *Frontiers in Human Neuroscience*, *9*, 12, Article 404. https://doi.org/10.3389/fnhum.2015.00404
Sales, A. C., Friston, K. J., Jones, M. W., Pickering, A. E., & Moran, R. J. (2019). Locus Coeruleus tracking of prediction errors optimises cognitive flexibility: An Active Inference model [Article]. *PLoS Computational Biology*, *15*(1), 24, Article e1006267. https://doi.org/10.1371/journal.pcbi.1006267
Schlaug, G. (2015). Musicians and music making as a model for the study of brain plasticity. In E. Altenmuller, S. Finger, & F. Boller (Eds.), *Music, Neurology, and Neuroscience: Evolution, the Musical Brain, Medical Conditions, and Therapies* (Vol. 217, pp. 37-55). Elsevier Science Bv. https://doi.org/10.1016/bs.pbr.2014.11.020
Schultz, W. (2015). Neuronal reward and decision signals: from theories to data [Review]. *Physiological Reviews*, *95*(3), 853-951. https://doi.org/10.1152/physrev.00023.2014
Schurr, R., Reznik, D., Hillman, H., Bhui, R., & Gershman, S. J. (2024). Dynamic computational phenotyping of human cognition [Article]. *Nature Human Behaviour*, *8*(5), 18. https://doi.org/10.1038/s41562-024-01814-x
Schutte, N. S., & Malouff, J. M. (2020). A Meta-Analysis of the Relationship between Curiosity and Creativity [Article]. *Journal of Creative Behavior*, *54*(4), 940-947. https://doi.org/10.1002/jocb.421
Schwartenbeck, P., & Friston, K. (2016). Computational Phenotyping in Psychiatry: A Worked Example [Article]. *Eneuro*, *3*(4), 18, Article 0049-16.2016. https://doi.org/10.1523/eneuro.0049-16.2016
Schwartenbeck, P., Passecker, J., Hauser, T. U., FitzGerald, T. H. B., Kronbichler, M., & Friston, K. J. (2019). Computational mechanisms of curiosity and goal-directed exploration [Article]. *Elife*, *8*, 45, Article e41703. https://doi.org/10.7554/eLife.41703
Schwarz, N. (2007). Attitude construction: Evaluation in context [Article]. *Social Cognition*, *25*(5), 638-656. https://doi.org/10.1521/soco.2007.25.5.638
Seger, C. A., & Spiering, B. J. (2011). A Critical Review of Habit Learning and the Basal Ganglia [Hypothesis and Theory]. *Frontiers in Systems Neuroscience*, *volume 5 - 2011*. https://doi.org/10.3389/fnsys.2011.00066
Sescousse, G., Caldú, X., Segura, B., & Dreher, J. C. (2013). Processing of primary and secondary rewards: A quantitative meta-analysis and review of human functional neuroimaging studies [Review]. *Neuroscience and Biobehavioral Reviews*, *37*(4), 681-696. https://doi.org/10.1016/j.neubiorev.2013.02.002
Shapiro, L. (2011). *Embodied Cognition*. Routledge. <Go to ISI>://WOS:000399995800009
Shettleworth, S. J. (2012). Do Animals Have Insight, and What Is Insight Anyway? [Article]. *Canadian Journal of Experimental Psychology-Revue Canadienne De Psychologie Experimentale*, *66*(4), 217-226. https://doi.org/10.1037/a0030674

Skov, M. (2019). Aesthetic Appreciation: The View From Neuroimaging [Article]. *Empirical Studies of the Arts*, *37*(2), 220-248. https://doi.org/10.1177/0276237419839257
Skov, M. (2023). Sensory liking: How nervous systems assign hedonic value to sensory objects. In M. S. M. Nadal (Ed.), *The Routledge international handbook of neuroaesthetics* (pp. 31-62). Routledge.
Skov, M., & Nadal, M. (2018). Art is not special: an assault on the last lines of defense against the naturalization of the human mind [Article]. *Reviews in the Neurosciences*, *29*(6), 699-702. https://doi.org/10.1515/revneuro-2017-0085
Skov, M., & Nadal, M. (2020). A Farewell to Art: Aesthetics as a Topic in Psychology and Neuroscience [Article]. *Perspectives on Psychological Science*, *15*(3), 630-642, Article 1745691619897963. https://doi.org/10.1177/1745691619897963
Smith, J. D., & Melara, R. J. (1990). AESTHETIC PREFERENCE AND SYNTACTIC PROTOTYPICALITY IN MUSIC - TIS THE GIFT TO BE SIMPLE [Article]. *Cognition*, *34*(3), 279-298. https://doi.org/10.1016/0010-0277(90)90007-7
Smith, R., Badcock, P., & Friston, K. J. (2021). Recent advances in the application of predictive coding and active inference models within clinical neuroscience [Article]. *Psychiatry and Clinical Neurosciences*, *75*(1), 3-13. https://doi.org/10.1111/pcn.13138
Smith, R., Friston, K. J., & Whyte, C. J. (2022). A step-by-step tutorial on active inference and its application to empirical data [Article]. *Journal of Mathematical Psychology*, *107*, 60, Article 102632. https://doi.org/10.1016/j.jmp.2021.102632
Snell-Rood, E. C. (2013). An overview of the evolutionary causes and consequences of behavioural plasticity [Article]. *Animal Behaviour*, *85*(5), 1004-1011. https://doi.org/10.1016/j.anbehav.2012.12.031
Spaan, M. T. J. (2012). Partially Observable Markov Decision Processes. In M. Wiering & M. van Otterlo (Eds.), *Reinforcement Learning: State-of-the-Art* (pp. 387-414). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-27645-3_12
Sterling, P. (2012). Allostasis: A model of predictive regulation [Article]. *Physiology and Behavior*, *106*(1), 5-15. https://doi.org/10.1016/j.physbeh.2011.06.004
Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction, 2nd Edition*. Mit Press. <Go to ISI>://WOS:000481873900019
Thoma, M. V., Ryf, S., Mohiyeddini, C., Ehlert, U., & Nater, U. M. (2012). Emotion regulation through listening to music in everyday situations [Article]. *Cognition & Emotion*, *26*(3), 550-560. https://doi.org/10.1080/02699931.2011.595390
Thomas, J. M., Higgs, S., Dourish, C. T., Hansen, P. C., Harmer, C. J., & McCabe, C. (2015). Satiation attenuates BOLD activity in brain regions involved in reward and increases activity in dorsolateral prefrontal cortex: an fMRI study in healthy volunteers [Article]. *American Journal of Clinical Nutrition*, *101*(4), 697-704. https://doi.org/10.3945/ajcn.114.097543
Thorndike, E. L. (1911). *Animal intelligence: Experimental studies* [doi:10.5962/bhl.title.55072]. Macmillan Press. https://doi.org/10.5962/bhl.title.55072
Topolinski, S., & Reber, R. (2010). Gaining Insight Into the "Aha" Experience [Article]. *Current Directions in Psychological Science*, *19*(6), 402-405. https://doi.org/10.1177/0963721410388803
Turk-Browne, N. B. (2012). Statistical Learning in Perception. In N. M. Seel (Ed.), *Encyclopedia of the Sciences of Learning* (pp. 3182-3185). Springer US. https://doi.org/10.1007/978-1-4419-1428-6_1707
Ullén, F., Hambrick, D. Z., & Mosing, M. A. (2016). Rethinking Expertise: A Multifactorial Gene-Environment Interaction Model of Expert Performance [Article]. *Psychological Bulletin*, *142*(4), 427-446. https://doi.org/10.1037/bul0000033
Unoki, S., Matsumoto, Y., & Mizunami, M. (2006). Roles of octopaminergic and dopaminergic neurons in mediating reward and punishment signals in insect visual learning [Article].

European Journal of Neuroscience, *24*(7), 2031-2038. https://doi.org/10.1111/j.1460-9568.2006.05099.x
Van de Cruys, S. (2017). Affective Value in the Predictive Mind. In T. K. Metzinger & W. Wiese (Eds.), *Philosophy and Predictive Processing*. MIND Group. https://doi.org/10.15502/9783958573253
Van de Cruys, S., Frascaroli, J., & Friston, K. (2024). Order and change in art: towards an active inference account of aesthetic experience [Review]. *Philosophical Transactions of the Royal Society B-Biological Sciences*, *379*(1895), 15, Article 20220411. https://doi.org/10.1098/rstb.2022.0411
van den Bosch, I., Salimpoor, V. N., & Zatorre, R. J. (2013). Familiarity mediates the relationship between emotional arousal and pleasure during music listening [Article]. *Frontiers in Human Neuroscience*, *7*, 10, Article 534. https://doi.org/10.3389/fnhum.2013.00534
von Helmholtz, H. (1866). *Concerning the perceptions in general* (Vol. 3). Dover.
Wiener, N. (1948). *Cybernetics; or control and communication in the animal and the machine*. John Wiley.
Winkielman, P., & Berridge, K. (2003). Irrational wanting and subrational liking: How rudimentary motivational and affective processes shape preferences and choices [Review]. *Political Psychology*, *24*(4), 657-680. https://doi.org/10.1046/j.1467-9221.2003.00346.x
Xiang, T., Ray, D., Lohrenz, T., Dayan, P., & Montague, P. R. (2012). Computational Phenotyping of Two-Person Interactions Reveals Differential Neural Response to Depth-of-Thought [Article]. *PLoS Computational Biology*, *8*(12), 9, Article e1002841. https://doi.org/10.1371/journal.pcbi.1002841
Yamamoto, K., & Vernier, P. (2011). The evolution of dopamine systems in chordates [Review]. *Frontiers in Neuroanatomy*, *5*, 21, Article 21. https://doi.org/10.3389/fnana.2011.00021
Zador, A., Escola, S., Richards, B., Ölveczky, B., Bengio, Y., Boahen, K., Botvinick, M., Chklovskii, D., Churchland, A., Clopath, C., DiCarlo, J., Ganguli, S., Hawkins, J., Körding, K., Koulakov, A., LeCun, Y., Lillicrap, T., Marblestone, A., Olshausen, B.,…Tsao, D. (2023). Catalyzing next-generation Artificial Intelligence through NeuroAI [Article]. *Nature Communications*, *14*(1), 7, Article 1597. https://doi.org/10.1038/s41467-023-37180-x
Zajonc, R. B. (1968). Attitudinal effects of mere exposure [Article]. *Journal of Personality and Social Psychology*, *9*(2P2), 1-&. https://doi.org/10.1037/h0025848
Zatorre, R. J. (2015). Musical pleasure and reward: mechanisms and dysfunction. In E. Bigand, B. Tillmann, I. Peretz, R. J. Zatorre, L. Lopez, & M. Majno (Eds.), *Neurosciences and Music V: Cognitive Stimulation and Rehabilitation* (Vol. 1337, pp. 202-211). Blackwell Science Publ. https://doi.org/10.1111/nyas.12677
Zatorre, R. J., Chen, J. L., & Penhune, V. B. (2007). When the brain plays music: auditory-motor interactions in music perception and production [Review]. *Nature Reviews Neuroscience*, *8*(7), 547-558. https://doi.org/10.1038/nrn2152
Zatorre, R. J., & Salimpoor, V. N. (2013). From perception to pleasure: Music and its neural substrates [Article]. *Proceedings of the National Academy of Sciences of the United States of America*, *110*, 10430-10437. https://doi.org/10.1073/pnas.1301228110