

REVIEW

The two kinds of free energy and the Bayesian revolution

Sebastian Gottwald ^{*}, Daniel A. Braun

Institute of Neural Information Processing, Ulm University, Ulm, Germany

^{*} sebastian.gottwald@uni-ulm.de

Abstract

The concept of free energy has its origins in 19th century thermodynamics, but has recently found its way into the behavioral and neural sciences, where it has been promoted for its wide applicability and has even been suggested as a fundamental principle of understanding intelligent behavior and brain function. We argue that there are essentially two different notions of free energy in current models of intelligent agency, that can both be considered as applications of Bayesian inference to the problem of action selection: one that appears when trading off accuracy and uncertainty based on a general maximum entropy principle, and one that formulates action selection in terms of minimizing an error measure that quantifies deviations of beliefs and policies from given reference models. The first approach provides a normative rule for action selection in the face of model uncertainty or when information processing capabilities are limited. The second approach directly aims to formulate the action selection problem as an inference problem in the context of Bayesian brain theories, also known as Active Inference in the literature. We elucidate the main ideas and discuss critical technical and conceptual issues revolving around these two notions of free energy that both claim to apply at all levels of decision-making, from the high-level deliberation of reasoning down to the low-level information processing of perception.


 OPEN ACCESS

Citation: Gottwald S, Braun DA (2020) The two kinds of free energy and the Bayesian revolution. *PLoS Comput Biol* 16(12): e1008420. <https://doi.org/10.1371/journal.pcbi.1008420>

Editor: Samuel J. Gershman, Harvard University, UNITED STATES

Published: December 3, 2020

Copyright: © 2020 Gottwald, Braun. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: This study was funded by the European Research Council under the Starting Grant ERC-StG-2015-ERC termed "BRISC: Bounded Rationality in Sensorimotor Coordination" with Project ID 678082 (D.A.B.). The first author's (S.G.) salary was paid from this funding. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

1 Introduction

There is a surprising line of thought connecting some of the greatest scientists of the last centuries, including Immanuel Kant, Hermann von Helmholtz, Ludwig E. Boltzmann, and Claude E. Shannon, whereby model-based processes of action, perception, and communication are explained with concepts borrowed from statistical physics. Inspired by Kant's Copernican revolution and motivated from his own studies of the physiology of the sensory system, Helmholtz was one of the first proponents of the *analysis-by-synthesis* approach to perception [1], whereby a perceiver is not simply conceptualized as some kind of tabula rasa recording raw external stimuli, but rather relies on internal models of the world to match and anticipate sensory inputs. The internal model paradigm is now ubiquitous in the cognitive and neural sciences and has even led some researchers to propose a Bayesian brain hypothesis, whereby the brain would essentially be a prediction and inference engine based on internal models [2–4]. Coincidentally, Helmholtz also invented the notion of the *Helmholtz free energy* that plays an

important role in thermodynamics and statistical mechanics, even though he never made a connection between the two concepts in his lifetime.

This connection was first made by Dayan, Hinton, Neal, and Zemel in their computational model of perceptual processing as a statistical inference engine known as the *Helmholtz machine* [5]. In this neural network architecture, there are feed-forward and feedback pathways, where the bottom-up pathway translates inputs from the bottom layer into hidden causes at the upper layer (the recognition model), and top-down activation translates simulated hidden causes into simulated inputs (the generative model). When considering log-likelihood in this setup as energy in analogy to statistical mechanics, learning becomes a relaxation process that can be described by the minimization of variational free energy. While it should be emphasized that variational free energy is not the same as Helmholtz free energy, the two free energy concepts can be formally related. Importantly, variational free energy minimization is not only a hallmark of the Helmholtz machine, but of a more general family of inference algorithms, such as the popular expectation-maximization (EM) algorithm [6, 7]. In fact, over the last two decades, variational Bayesian methods have become one of the foremost approximation schemes for tractable inference in the machine learning literature. Moreover, a plethora of machine learning approaches use loss functions that have the shape of a free energy when optimizing performance under entropy regularization in order to boost generalization of learning models [8, 9].

In the meanwhile, free energy concepts have also made their way into the behavioral sciences. In the economic literature, for example, trade-offs between utility and entropic uncertainty measures that take the form of free energies have been proposed to describe decision-makers with stochastic choice behavior due to limited resources [10–14] or robust decision-makers with limited precision in their models [15, 16]. The free energy trade-off between entropy and reward can also be found in information-theoretic models of biological perception-action systems [17–19], some of which have been subjected to experimental testing [20–25]. Finally, in the neuroscience literature the notion of free energy has risen to recent fame as the central puzzle piece in the Free Energy Principle [26] that has been used to explain a cornucopia of experimental findings including neural prediction error signals [27], synaptic plasticity rules [28], neural effects of biased competition and attention [29, 30], visual exploration in humans [31], and more—see the references in [32]. Over time, the Free Energy Principle has grown out of an application of the free energy concept used in the Helmholtz machine, to interpret cortical responses in the context of predictive coding [33], and has gradually developed into a general principle for intelligent agency, also known as *Active Inference* [32, 34, 35]. Consequences and implications of the Free Energy Principle are discussed in neighbouring fields like psychiatry [36, 37] and the philosophy of mind [38, 39].

Given that the notion of free energy has become such a pervasive concept that cuts through multiple disciplines, the main rationale for this discussion paper is to trace back and to clarify different notions of free energy, to see how they are related and what role they play in explaining behavior and neural activity. As the notion of free energy mainly appears in the context of statistical models of cognition, the language of probabilistic models constitutes a common framework in the following discussion. Section 2 therefore starts with preliminary remarks on probabilistic modelling. Section 3 introduces two notions of free energy that are subsequently expounded in Section 4 and Section 5, where they are applied to models of intelligent agency. Section 6 concludes the paper.

2 Probabilistic models and perception-action systems

Systems that show stochastic behavior, for example due to randomly behaving components or because the observer ignores certain degrees of freedom, are modelled using probability

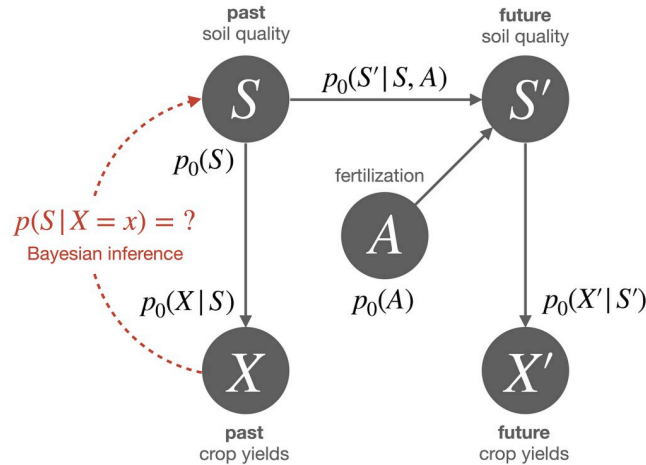


Fig 1. Graphical representation of an exemplary probabilistic model. The arrows (edges) indicate causal relationships between the random variables (nodes). The full joint distribution p_0 over all random variables is sometimes also referred to as a *generative model*, because it contains the complete knowledge about the random variables and their dependencies and therefore allows to generate simulated data. Such a model could for example be used by a farmer to infer the soil quality S based on the crop yields X through Bayesian inference, which allows to determine a priori unknown distributions such as $p(S|X)$ from the generative model p_0 via marginalization and conditionalization.

<https://doi.org/10.1371/journal.pcbi.1008420.g001>

distributions. This way, any behavioral, environmental, and hidden variables can be related by their statistics, and dynamical changes can be modelled by changes in their distributions.

Consider, for example, the simple probabilistic model illustrated in Fig 1, consisting of the (for simplicity, discrete) variables *past and future soil quality* $\mathbf{S} := (S, S')$, *past and future crop yields* $\mathbf{X} := (X, X')$, and *fertilization* A . The graphical model shown in the figure corresponds to the joint probability $p_0(\mathbf{X}, \mathbf{S}, A)$ given by the factorization

$$p_0(X'|S') p_0(X|S) p_0(S'|S, A) p_0(S) p_0(A), \tag{1}$$

where $p_0(S)$ is the base probability of the past soil quality S , $p_0(X|S)$ is the probability of crop yields X depending on the past soil quality S , and so forth. Given the joint distribution we can now ask questions about each of the variables. For example, we could ask about the probability distribution $p(S|X = x)$ of soil quality S if we are told that the crop yields X are equal to a value x . We can obtain the answer from the probabilistic model p_0 by doing *Bayesian inference*, yielding the *Bayes' posterior*

$$p(S|X) = \frac{p(S, X)}{\sum_s p(s, X)} = \frac{p_0(X|S)p_0(S)}{\sum_s p_0(X|s)p_0(s)}, \tag{2}$$

where the dependencies on X' , S' , and A have been summed out to calculate the marginal $p(S, X)$. In general, Bayesian inference in a probabilistic model means to determine the probability of some queried unobserved variables given the knowledge of some observed variables. This can be viewed as transforming the *prior* probabilistic model p_0 to a *posterior* model p , under which the observed values have probability one, and unobserved variables have probabilities given by the corresponding Bayes' posteriors.

In principle, Bayesian inference requires only two different kinds of operations, namely *marginalization*, i.e., summing out unobserved variables that have not been queried, such as X' , S' and A above, and *conditionalization*, i.e., renormalizing the joint distribution over observed and queried variables—that may itself be the result from a previous marginalization

such as $p(S, X)$ above—to obtain the required conditional distribution over the queried variables. In practice, however, inference is a hard computational problem and many more efficient inference methods are available that may provide approximate solutions to the exact Bayes' posteriors, including belief propagation [40], expectation propagation [41], variational Bayesian inference [42], and Monte Carlo algorithms [43]. Also note that inference is trivial if the sought-after conditional distribution of the queried variable is already given by one of the conditional distributions that jointly specify the probabilistic model, e.g., $p(X|S) = p_0(X|S)$.

Probabilistic models can be used not only as *external (observer) models*, but also as *internal models* that are employed by the agent itself, or by a designer of the agent, in order to determine a desired course of action. In this latter case, actions could either be thought of as deterministic parameters of the probabilistic model that influence the future (*influence diagrams*) or as random variables that are part of the probabilistic model themselves (*prior models*) [44]. Either way, internal models allow making predictions over future consequences in order to find actions or distributions over actions that lead to desirable outcomes, for example actions that produce high rewards in the future. In *mechanistic* or *process* model interpretations, some of the specification procedures to find such actions are themselves meant to represent what the agent is actually doing while reasoning, whereas *as if* interpretations simply use these methods as tools to arrive at distributions that describe the agent's behavior. Free energy is one of the concepts that appears in both types of methods.

3 The two notions of free energy

Vaguely speaking, free energy can refer to any quantity that is of the form

$$\text{free energy} = \text{energy} \pm \text{const.} \times \text{entropy}, \quad (3)$$

where *energy* is an expected value of some quantity of interest, *entropy* refers to a quantity measuring disorder, uncertainty, or complexity, that must be specified in the given context, and *const.* is a constant term that translates between units of entropy and energy, and is related to the *temperature* in physically motivated free energy expressions. From relation (3), it is not surprising that free energy sometimes appears enshrouded by mystery, as it relies on an understanding of entropy, and “nobody really knows what entropy is anyway”, as John Von Neumann famously quipped [45].

Historically, the concept of free energy goes back to the roots of thermodynamics, where it was introduced to measure the maximum amount of work that can be extracted from a thermodynamic system at a constant temperature and volume. If, for example, all the molecules in a box move to the left, we can use this kinetic energy to drive a turbine. If, however, the same kinetic energy is distributed as random molecular motion, it cannot be fully transformed into work. Therefore, only part of the total energy E is usable, because the exact positions and momenta of the molecules, the so-called *microstates*, are unknown. In this case, the maximum usable part of the energy E is the *Helmholtz free energy*, defined as

$$F = E - TS, \quad (4)$$

where S is the thermodynamic entropy. In general, the transformation between two macrostates with free energies F_1 and F_2 allows the extraction of work $W \leq F_2 - F_1$.

While the two notions of free energy that we discuss in the following are vaguely inspired by the physical original, their motivations are rather distinct and the main reason they share the nomenclature is due to their general form (3) resembling the Helmholtz free energy (4).

3.1 Free energy from constraints

The first notion of free energy is closely tied to the *principle of maximum entropy* [46], which virtually appears in all branches of science. From this vantage point, the physical free energy is merely a special instance of a more general inference problem where we hold probabilistic beliefs about unknown quantities (e.g., the exact energy values of the molecules in a gas) and we can only make coarse measurements or observations (e.g., the temperature of the gas) that we can use to update our beliefs about these hidden variables. The principle of maximum entropy suggests that, among the beliefs that are compatible with the observation, we should choose the most “unbiased” belief, in the sense that it corresponds to a maximum number of possible assignments of the hidden variables.

3.1.1 Wallis’ motivation of the maximum entropy principle. Consider the random experiment of distributing N elements randomly in n equally probable buckets with $N \gg n$, where the resulting number of elements N_i in bucket $i \in \{1, \dots, n\}$ determines the probability $p(z_i) := \frac{N_i}{N}$. In principle, this way we could generate any distribution p over a finite set $\Omega = \{z_1, \dots, z_n\}$ that we like, however, a uniform distribution that reflects the equiprobable assignment clearly is much more likely than a Dirac distribution where all the probability mass is concentrated in one bucket. Here, the reason is that there are many possible assignments of elements among the buckets that generate the uniform distribution, whereas there is only one for a Dirac distribution. In fact, the number of possibilities of how to distribute N elements among n buckets with N_i elements in the i th bucket is

$$\omega := \frac{N!}{N_1! \dots N_n!}, \tag{5}$$

because $N!$ is the number of possible permutations of all N elements, which overcounts by the number of permutations of elements inside the same bucket and thus has to be divided by the number of permutations $N_i!$ for all $i = 1, \dots, n$. In the absence of any further measurement constraints, the number of possibilities (5) is maximized by $N_i = N/n$ for all i , and thus the *typical* distribution p^* over Ω in this case is the uniform distribution, i.e., $p^*(z_i) = \frac{1}{n}$ for all i .

Consider now the problem of having to determine a typical distribution p^* over Ω such that the expected value $\mathbb{E}_{p^*}[\mathcal{E}] = \langle \mathcal{E} \rangle_{p^*}$ of some quantity \mathcal{E} equals a measured value ε . A simple example would be the experiment of throwing N dice and taking \mathcal{E} to be the number of dots, i.e., $\mathcal{E}(z_1) = 1, \dots, \mathcal{E}(z_6) = 6$, and trying to find the typical distribution p^* over outcomes z_1, \dots, z_6 under the constraint that the average number of dots is, say $\varepsilon = 2$. The solution to this problem is analogous to the case of no constraints, but this time we only consider realizations that are compatible with the measurement constraint, that is we let (N_1, \dots, N_n) belong to the set of permissible occupation vectors

$$\Gamma_\varepsilon := \{(N_1, \dots, N_n) \mid \langle \mathcal{E} \rangle_p = \varepsilon, p(x_i) = \frac{N_i}{N} \forall i\}.$$

A typical distribution p^* for a constraint ε can then be determined by a candidate in Γ_ε with the maximum number ω of possibilities (5). By assumption, N is much larger than n , so that we can get rid of the faculties by making use of Stirling’s approximation $\ln N! = N \ln N - N + \mathcal{O}(\ln N)$. In particular, when letting $N, N_i \rightarrow \infty$ such that $p(x_i) = \frac{N_i}{N}$ remains finite, we obtain

$$\frac{1}{N} \log \omega = \underbrace{-\sum_{i=1}^n \frac{N_i}{N} \log \frac{N_i}{N}}_{=H(p)} + \mathcal{O}\left(\frac{\log N}{N}\right) \xrightarrow{N \rightarrow \infty} H(p).$$

where $H(p) := -\sum_{z \in \Omega} p(z) \log p(z)$ denotes the (*Gibbs or Shannon*) *entropy* of p . Thus, instead of assessing typicality by maximizing (5) in Γ_ϵ for large but fixed N , we can get rid of the N -dependency by simply maximizing H ,

$$p^* = \operatorname{argmax}_{p, \langle \mathcal{E} \rangle_p = \epsilon} H(p). \tag{6}$$

This constrained optimization problem is known as the principle of maximum entropy. The motivation given here is essentially the Wallis derivation presented by Jaynes [47].

3.1.2 Free energy from constraints and the Boltzmann distribution. The constrained optimization problem (6) can be translated into an unconstrained problem by introducing a Lagrange multiplier β , known as the *inverse temperature* due to the analogy to thermodynamics and the Helmholtz Free Energy (4), which has to be chosen post hoc such that the constraint is satisfied. This results in the minimization of the Lagrangian

$$F(p) := \langle \mathcal{E} \rangle_p - \frac{1}{\beta} H(p), \tag{7}$$

which takes the form of a free energy (3). As we shall see later, F takes its minimum at the *Boltzmann distribution* known from statistical mechanics, given by

$$p^*(z) := \frac{1}{\mathcal{Z}} e^{-\beta \mathcal{E}(z)}, \tag{8}$$

where $\mathcal{Z} = \sum_{z \in \Omega} e^{-\beta \mathcal{E}(z)}$ denotes the normalization constant.

Note that, the argument in the previous section implicitly assumes a *uniform* reference distribution, because the buckets are assumed to be equiprobable. When replacing this assumption by the assumption of a general distribution p_0 over Ω , we obtain the *principle of minimum relative entropy* [48], where the so-called *Kullback-Leibler (KL) divergence* $D_{\text{KL}}(p \| p_0) = \langle \log(p/p_0) \rangle_p$ is minimized with respect to p subject to a constraint $\langle \mathcal{E} \rangle_p = \epsilon$. Analogous to the maximum entropy principle, this translates to the unconstrained minimization of the Lagrangian

$$F(p, p_0) := \langle \mathcal{E} \rangle_p + \frac{1}{\beta} D_{\text{KL}}(p \| p_0), \tag{9}$$

with solution given by $p^*(z) = \frac{1}{\mathcal{Z}} p_0(z) e^{-\beta \mathcal{E}(z)}$.

3.1.3 The trade-off between energy and uncertainty. An important feature of the minimization of the free energies (7) and (9) consists in the balancing of the two competing terms of energy and entropy (cf. Fig 2). This *trade-off* between maximal uncertainty (uniform distribution, or p_0) on the one hand and minimal energy (e.g., a delta distribution) on the other hand is the core of the maximum entropy principle. The inverse temperature β plays the role of a trade-off parameter that controls how these two counteracting forces are weighted.

The maximum entropy principle goes back to the *principle of insufficient reason* [49–51], which states that two events should be assigned the same probability if there is no reason to think otherwise. It has been hailed as a principled method to determine prior distributions and to incorporate novel information into existing probabilistic knowledge. In fact, Bayesian inference can be cast in terms of relative entropy minimization with constraints given by the available information [52]. Applications of this idea can also be found in the machine learning literature, where subtracting (or adding) an entropy term from an expected value of a function that must be optimized is known as *entropy regularization* and plays an important role in modern reinforcement learning algorithms [8, 9] to encourage exploration [53] as well as to penalize overly deterministic policies resulting in biased reward estimates [54].

From now on, we refer to a free energy expression that is motivated from a trade-off between an energy and an entropy term, such as (7) and (9), as *free energy from constraints*, in




Fig 2. Minimizing the free energy from constraints (7) requires to trade off the competing terms of energy $\langle \mathcal{E} \rangle_p$ and entropy $H(p)$, here shown exemplarily for the case of three elements. Assuming there exists a unique minimal element $z^* = \operatorname{argmin}_z \mathcal{E}(z)$, then minimizing only $\langle \mathcal{E} \rangle_p$ over all probability distributions p results in the (Dirac delta) distribution δ_{z^*} that assigns zero probability to all $z_i \neq z^*$ and probability one to $z_i = z^*$, and therefore has zero entropy. In contrast, minimizing only the term $-\frac{1}{\beta} H(p)$ is equivalent to maximizing $H(p)$ and therefore would result in the uniform distribution that gives equal probability to all elements. The resulting Boltzmann distribution p^* interpolates between these two extreme solutions of minimal energy ($\beta \rightarrow \infty$) and maximum entropy ($\beta \rightarrow 0$).

<https://doi.org/10.1371/journal.pcbi.1008420.g002>

order to discriminate it from the notion of free energy introduced in the following section, which—despite of its resemblance—has a different motivation.

3.2 Variational free energy

There is another, distinct appearance of the term “free energy” outside of physics, which is a priori *not* motivated from a trade-off between an energy and entropy term, but from possible efficiency gains when representing Bayes’ rule in terms of an optimization problem. This technique is mainly used in *variational Bayesian inference* [55], originally introduced by Hinton and van Camp [42]. As before, for simplicity all random variables are discrete, but most expressions can directly be translated to the continuous case by replacing probability distributions by probability densities and sums by the corresponding integrals.

3.2.1 Variational Bayesian inference. As we have seen in Section 2, Bayesian inference consists in the calculation of a conditional probability distribution over unknown variables given the values of known variables. In the most simple case of two variables, say X and Z , and a probabilistic model of the form $p_0(X, Z) = p_0(X|Z)p_0(Z)$, Bayesian inference applies if X is observed and Z is queried. Analogous to (2), the *exact* Bayes’ posterior $p(Z|X = x)$ is defined by the renormalization of $p_0(x, Z)$ in order to obtain a distribution over Z that respects the new information $X = x$,

$$p(Z|X = x) = \frac{p_0(x, Z)}{\mathcal{Z}(x)} = \frac{p_0(x|Z)p_0(Z)}{\mathcal{Z}(x)}, \tag{10}$$

with the normalization constant $\mathcal{Z}(x) = \sum_z p_0(x, z) = p(X = x)$.

In *variational Bayesian inference*, however, this Bayes’ posterior is not calculated directly by renormalizing $p_0(x, Z)$ with respect to Z , but indirectly by approximating it by a distribution $q(Z)$ that is adjusted through the minimization of an error measure that quantifies the deviation from the exact Bayes’ posterior. Importantly, the value of this error measure can be determined without having to know the exact Bayes’ posterior. To see this, note that the KL

divergence between $q(Z)$ and $p(Z|X = x)$ can be written as

$$\underbrace{\left\langle \log \frac{q(Z)}{p(Z|X = x)} \right\rangle_{q(Z)}}_{=D_{\text{KL}}(q(Z)\|p(Z|X=x))} = \underbrace{\log \mathcal{Z}(x)}_{\text{indep. of } q} + \underbrace{\left\langle \log \frac{q(Z)}{p_0(x, Z)} \right\rangle_{q(Z)}}_{=:F(q(Z)\|p_0(x, Z))}, \tag{11}$$

i.e., it can be decomposed into the sum of a constant term and a term that does not depend on the normalization $\mathcal{Z}(x)$. In particular, a good approximation $q(Z)$ of the exact Bayes’ posterior (10) will effectively minimize this KL divergence, which—due to (11)—can be done by minimizing $F(q(Z)\|p_0(x, Z))$. In particular, the optimum of this minimization is exactly achieved at the Bayes’ posterior (10),

$$\operatorname{argmin}_{q(Z)} \left\langle \log \frac{q(Z)}{p_0(x, Z)} \right\rangle_{q(Z)} = p(Z|X = x), \tag{12}$$

which is known as the variational characterization of Bayes’ rule. This result is a special case of (14) in the following section.

3.2.2 Variational free energy, an extension of relative entropy. Any non-negative function ϕ on a finite space Ω , can be normalized to obtain a probability distribution $p_\phi = \phi/\sum_z \phi(z)$ on Ω that differs from ϕ only by a scaling constant. In cases when it is not beneficial to carry out the sum $\sum_z \phi(z)$ explicitly, such a normalization might be replaced by the minimization of the *variational free energy*

$$F(q\|\phi) := \left\langle \log \frac{q(Z)}{\phi(Z)} \right\rangle_{q(Z)}, \tag{13}$$

with respect to the so-called *trial distributions* q , because we have

$$\operatorname{argmin}_q F(q\|\phi) = \frac{\phi(Z)}{\sum_z \phi(z)} = p_\phi(Z). \tag{14}$$

Thus, instead of normalizing ϕ directly, one fits auxiliary distributions q to approximate the shape of ϕ in the space of probability distributions (cf. Fig 3). If this optimization process has no constraints, then the trial distributions are adjusted until p_ϕ is achieved. In the case of constraints, for instance if the trial distributions are parametrized by a non-exhaustive parametrization (e.g., Gaussians), then the optimized trial distributions approximate p_ϕ as close as possible within this parametrization. The minimal value of $F(q\|\phi)$ is

$$F(p_\phi\|\phi) = \min_q F(q\|\phi) = -\log \sum_z \phi(z). \tag{15}$$

In particular, this implies that $-F(q\|\phi) \leq \log \sum_z \phi(z)$ for all q , so that varying $-F(q\|\phi)$ with arbitrary trial distributions q always provides a lower bound to the unknown normalization constant $\sum_z \phi(z)$. In Bayesian inference this is the normalization constant in Bayes’ rule and called the *model evidence*, which is why the negative variational free energy is also called *evidence lower bound* (ELBO).

The proof of (14) and (15) directly follows from Jensen’s inequality and only relies on the concavity of the logarithm. As we have seen in the previous section, in variational Bayesian inference, the *reference* ϕ usually takes the form of a joint distribution evaluated at the observed variables, e.g., $\phi(Z) = p_0(x, Z)$ in which case (14) recovers (12). The variational free energy (13) is a free energy in the sense of (3) since by the additivity of the logarithm under

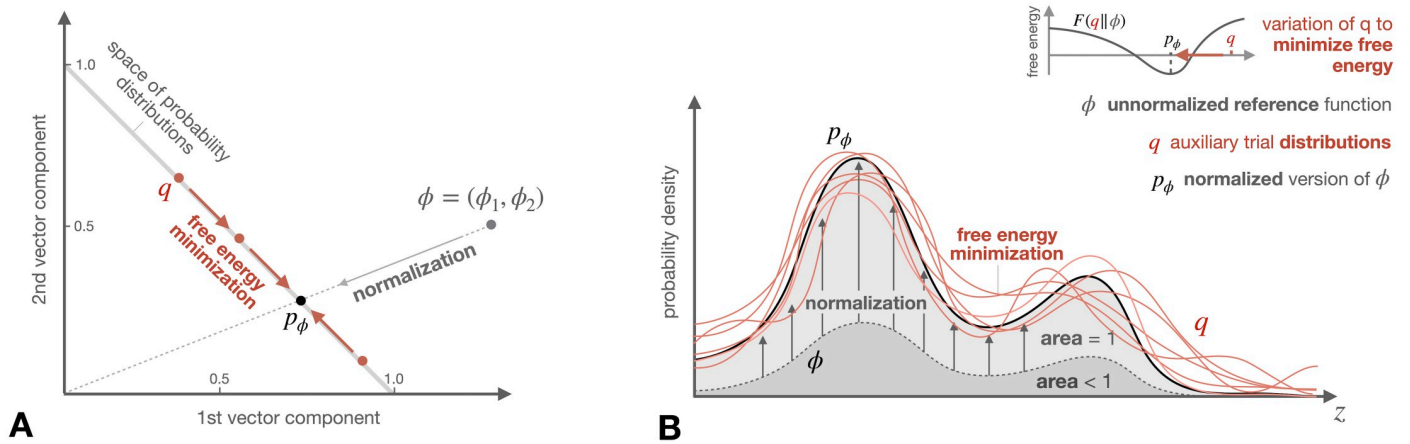


Fig 3. The normalization of a function ϕ to obtain a probability distribution p_ϕ is equivalent to fitting trial distributions q to the shape of ϕ by minimizing free energy. In two dimensions, the normalization of a point $\phi = (\phi_1, \phi_2)$ corresponds to a (non-orthogonal) projection onto the plane of probability vectors (A). For continuous domains, where probability distributions are represented by densities, normalization corresponds to a rescaling of ϕ such that the area below the graph equals 1 (B). Instead, when minimizing variational free energy (red colour), the trial distributions q are varied until they fit to the shape of the unnormalized function ϕ (perfectly at $q = p_\phi$).

<https://doi.org/10.1371/journal.pcbi.1008420.g003>

multiplication ($\log ab = \log a + \log b$),

$$F(q||\phi) = \langle -\log\phi \rangle_q - H(q) \tag{16}$$

with energy term $\langle -\log\phi \rangle_q$ and entropy term $H(q)$. Note that, for the choice $\phi = e^{-\beta\epsilon}$, Eq (14) becomes the Boltzmann distribution (8) and the variational free energy (16) formally corresponds to the free energy from constraints (7).

Variational free energy can be regarded as an extension of relative entropy with the reference distribution being replaced by a non-normalized reference function, since in the case when ϕ is already normalized, that is if $\sum_z \phi(z) = 1$, then the free energy (13) coincides with the KL divergence $D_{\text{KL}}(q||\phi)$. In particular, while relative entropy is a measure for the dissimilarity of two probability distributions, where the minimum is achieved if both distributions are equal, variational free energy is a measure for the dissimilarity between a probability distribution q and a (generally non-normalized) function ϕ , where the minimum with respect to q is achieved at p_ϕ . Accordingly, we can think of the variational free energy as a specific error measure between probability distributions and reference functions. In principle, one could design many other error measures that have the same minimum. This means that, a statement in a probabilistic setting that a distribution q^* minimizes a variational free energy $F(q||\phi)$ with respect to a given reference ϕ , is analogous to a statement in a non-probabilistic setting that some number $x = x^*$ minimizes the value of an error measure $\epsilon(x, y)$ (e.g., the squared error $\epsilon(x, y) = (x - y)^2$) with respect to a given reference value y .

3.2.3 Approximate and iterative inference. Representing Bayes' rule as an optimization problem over auxiliary distributions q has two main applications that both can simplify the inference process (cf. Fig 4). First, it allows to *approximate* exact Bayes' posteriors by restricting the optimization space, for example using a non-exhaustive parametrization, e.g., an exponential family. Second, it enables *iterative inference algorithms* consisting of multiple simpler optimization steps, for example by optimizing with respect to each term in a factorized representation of q separately. A popular choice is the *mean-field approximation*, which combines both of these simplifications, as it assumes independence between hidden states, effectively reducing the search space from joint distributions to factorized ones, and moreover it allows

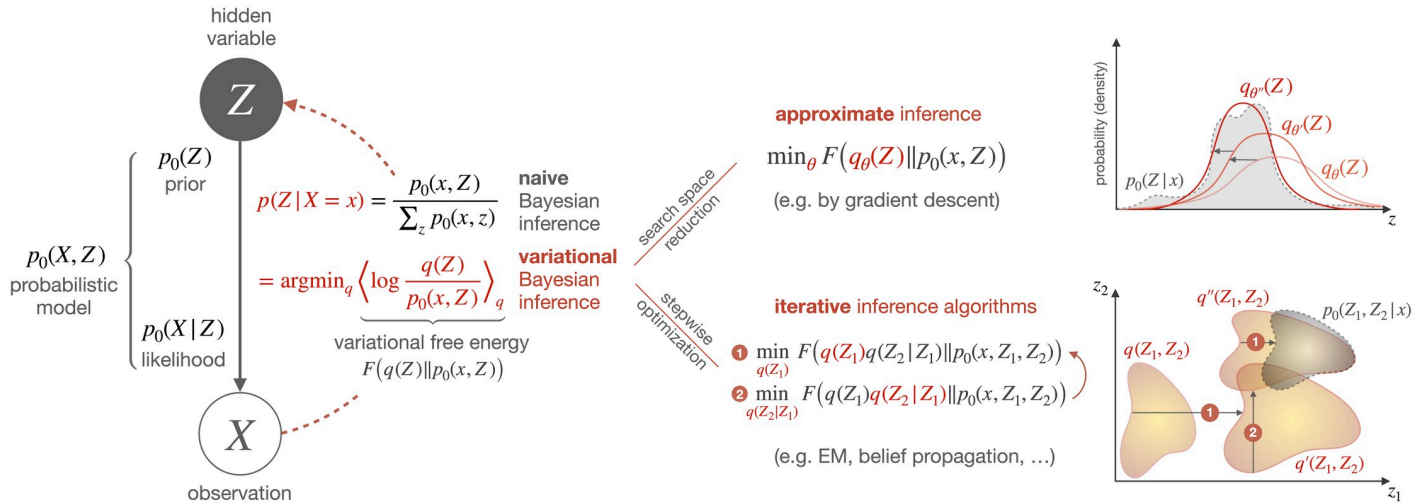


Fig 4. In variational Bayesian inference, the operation of renormalizing the probabilistic model p_0 evaluated at an observation $X = x$ (Bayes' rule), is replaced by an optimization problem. In practice, this variational representation is often exploited to simplify a given inference problem, either by reducing the search space of distributions, for example through a restrictive parametrization resulting in approximate inference, or by splitting up the optimization into multiple partial optimization steps that are potentially easier to solve than the original problem but might still converge to the exact solution. These two simplifications can also be combined, for example in the case of mean-field assumptions where the space of distributions is reduced *and* an efficient iterative inference algorithm is obtained at the same time.

<https://doi.org/10.1371/journal.pcbi.1008420.g004>

to optimize with respect to each factor alternately. Note, however, that mean-field approximations have limited use in sequential environments, where independence of subsequent states cannot be assumed and therefore less restrictive assumptions must be used instead [56].

Many efficient iterative algorithms for exact and approximate inference can be viewed as examples of variational free energy minimization, for example the EM algorithm [6, 57], belief propagation [40, 58], and other message passing algorithms [41, 59–62]. While the (Bayesian) EM algorithm [7] and Pearl's belief propagation [58] both can be seen as minimizing the same variational free energy, just with different assumptions on the approximate posteriors, in [61], it is shown that also many other message passing algorithms such as [41, 59, 60] can be cast as minimizing some type of free energy, the only difference being the choice of the divergence measure as the entropy term. Simple versions of these algorithms have often existed before their free energy formulations were available, but the variational representations usually allowed for extensions and refinements—see [6, 7, 63, 64] in case of EM and [58, 62, 65, 66] in case of message passing.

We are now turning to the question of how the two notions of free energy introduced in this section are related to recent theories of intelligent agency.

4 Free energy from constraints in information processing

4.1 The basic idea

The concept of free energy from constraints as a trade-off between energy and uncertainty can be used in models of perception-action systems, where entropy quantifies information processing complexity required for decision-making (e.g., planning a path for fleeing a predator) and energy corresponds to performance (e.g., distinguishing better and worse flight directions). The notion of decision in this context is very broad and can be applied to any internal variable in the perception-action pipeline [67], that is not given directly by the environment.

In particular, it also subsumes perception itself, where the decision variables are given by the hidden causes that are being inferred from observations.

In rational choice theory [68], a decision-maker selects decisions x^* from a set of options Ω such that a utility function U defined on Ω is maximized,

$$x^* = \operatorname{argmax}_{x \in \Omega} U(x). \quad (17)$$

The utility values $U(x)$ could either be objective, for example a monetary gain, or subjective in which case they represent the decision-maker's preferences. In general, the utility does not have to be defined directly on Ω , but could be derived from utility values that are attached to certain states, for example to the configurations of the playboard in a board game. In the case of perception, utility values are usually given by (log-)likelihood functions, in which case utility maximization without constraints corresponds to greedy inference such as maximum likelihood estimation. Note that, for simplicity, in this section we consider one-step decision problems. Sequential tasks can either be seen as multiple one-step problems where the utility of a given step might depend on the policy over future steps, or as path planning problems where an action represents a full action path or policy [18, 69–71].

While ideal rational decision-makers are assumed to perfectly optimize a given utility function U , real behavior is often stochastic, meaning that multiple exposures to the same problem lead to different decisions. Such non-deterministic behavior could be a consequence of model uncertainty, as in Bayesian inference or various stochastic gambling schemes, or a consequence of *satisficing* [72], where decision-makers do not choose the single best option, but simply one option that is good enough. Abstractly, this means that, the choice of a single decision is replaced by the choice of a distribution over decisions. More generally, also considering prior information that the decision-maker might have from previous experience, the process of deliberation during decision-making might be expressed as the transformation of a prior p_0 to a posterior distribution p .

When assuming that deliberation has a cost $C(p, p_0)$, then arriving at narrow posterior distributions should intuitively be more costly than choosing distributions that contain more uncertainty (cf. Fig 5A). In other words, deliberation costs must be increasing with the amount of uncertainty that is reduced by the transformation from p_0 to p . Uncertainty reduction can be understood as making the probabilities of options less equal to each other, rigorously expressed by the mathematical concept of majorization [73]. This notion of uncertainty can also be generalized to include prior information, so that the degree of uncertainty reduction corresponds to more or less deviations from the prior [74].

Maximizing expected utility $\langle U \rangle_p$ with respect to p under restrictions on processing costs $C(p, p_0)$ is a constrained optimization problem that can be interpreted as a particular model of *bounded rationality* [72], explaining non-rational behavior of decision-makers that may be unable to select the single best option by their limited information processing capability. Similarly to the free energy trade-off between energy and entropy (cf. Fig 2), this results in a trade-off between utility $\langle U \rangle_p$ and processing costs $C(p, p_0)$,

$$F_\beta(p) := \langle U \rangle_p - \frac{1}{\beta} C(p, p_0). \quad (18)$$

Here, the trade-off parameter β is analogous to the *inverse temperature* in statistical mechanics (cf. Eq (7)) and parametrizes the optimal trade-offs $p_\beta^* = \operatorname{argmax}_p F_\beta(p)$ between utility and cost, that define an efficiency frontier separating the space of perception-action systems into bounded-optimal, non-optimal, and non-admissible systems (cf. Fig 5).




Fig 5. A: Decision-making can be considered as a search process in the space of options Ω , where options are progressively ruled out. Deliberation costs are defined to be monotone functions under such uncertainty reduction. **B:** Exemplary efficiency curve resulting from the trade-off between utility and costs, that separates non-optimal from non-admissible behavior. The points on the curve correspond to bounded-optimal agents that optimally trade off utility against uncertainty, analogous to the rate-distortion curve in information theory.

<https://doi.org/10.1371/journal.pcbi.1008420.g005>

When assuming that the total transformation cost is the same independent of whether a decision problem is solved in one step or multiple sub-steps (*additivity under coarse-graining*) the trade-off in (18) takes the general form (3) of a free energy in the sense of energy (utility) minus entropy (cost), because then the cost function is uniquely given by the relative entropy

$$C(p, p_0) = D_{\text{KL}}(p||p_0). \tag{19}$$

Note that the additivity of (19) also implies a coarse-graining property of the free energy (18) in the case when the decision is split into multiple steps, such that the utility of preceding decisions is effectively given by the free energy of following decisions. Therefore, in this case, free energy can be seen as a *certainty-equivalent* value of the subordinate decision problems, i.e., the amount of utility the agent would have to receive to be indifferent between this guaranteed utility and the potential expected utility of the subsequent decision steps taking account the associated information processing costs. The special case (19) has been studied extensively in multiple contexts, including quantal response equilibria in the game-theoretic literature [10, 14], rational inattention and costly contemplation [11, 75], bounded rationality with KL costs [12, 19], KL control [76, 77], entropy regularization [8, 9], robustness [15, 16], the emergence of heuristics [78], thermodynamic models of computation [79], and the analysis of information flow in perception-action systems [17, 18]. While (19) is often regarded as an abstract measure of uncertainty reduction or a generic proxy for information processing costs, it can also be viewed as a physical capacity constraint, where the information that is required to achieve a certain expected utility is considered to be sent over a channel to the actuator [24, 80–83]. This view is also consistent with the maximum entropy principle, as (18) and (19) favor distributions p that can be generated from p_0 most easily in terms of statistics, and therefore with minimum communication complexity between p_0 and p [84].

4.2 A simple example

Ingredients. Consider the probabilistic model shown in Fig 1 with the joint distribution $p_0(X, S, A)$ that is specified by the factors in the decomposition (1). Here, S and X denote the current

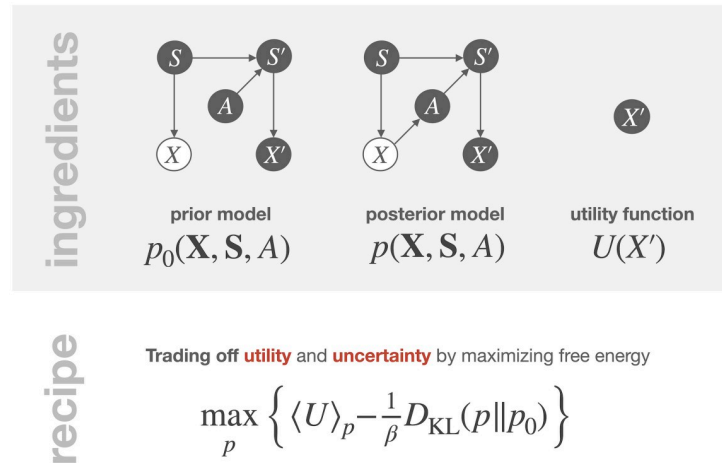


Fig 6. Overview of how to apply utility maximization with information processing costs to the example from Section 2.

<https://doi.org/10.1371/journal.pcbi.1008420.g006>

environmental state and the corresponding observation, and A denotes the action that must be determined in order to drive the system into a new state S' with observation X' . The decision-making problem is specified by assuming that we have given a utility function U over future observations X' which the decision-maker seeks to maximize by selecting an action A , while only having access to the current observation X . This means that the decision-maker has control over the distribution $p(A|X)$, which replaces the prior $p_0(A)$ in the factorization (1) of the prior model $p_0(\mathbf{X}, \mathbf{S}, A)$ to determine the factorization of the posterior model $p(\mathbf{X}, \mathbf{S}, A)$ in terms of the fixed components in p_0 (cf. Fig 6) as

$$p(\mathbf{X}, \mathbf{S}, A) = \underbrace{p_0(X'|S') p_0(X|S) p_0(S'|S, A) p_0(S)}_{p_0(\mathbf{X}, \mathbf{S}|A)} p(A|X). \tag{20}$$

Free energy from constraints. Further assuming that the decision-maker is subject to an information processing constraint $D_{\text{KL}}(p||p_0) \leq C_0$, for some non-negative bound C_0 , results in the unconstrained optimization problem $\max_p F(p)$ with free energy given by (18), where the trade-off parameter β is tuned to comply with the bound C_0 . Since the action distribution $p(A|X)$ is the only distribution in the posterior model (20) that changes during decision-making, i.e., during the transformation from prior to posterior, the total free energy simplifies to

$$\begin{aligned} F(p) &= \langle U \rangle_{p(\mathbf{X}, \mathbf{S}, A)} - \frac{1}{\beta} D_{\text{KL}}(p(\mathbf{X}, \mathbf{S}, A) || p_0(\mathbf{X}, \mathbf{S}, A)) \\ &= \langle V(X, A) \rangle_{p(A|X)p(X)} - \frac{1}{\beta} \langle D_{\text{KL}}(p(A|X) || p_0(A)) \rangle_{p(X)} \\ &= \langle F_A(p(A|X)) \rangle_{p(X)}, \end{aligned}$$

where we have written $p_0(x|s)p_0(s) = p(s|x)p(x)$ using Bayes' rule (2), and

$$\begin{aligned} V(X, A) &:= \sum_{s, s', x'} p(s|X) p_0(s'|s, A) p_0(x'|s') U(x'), \\ F_A(p(A|X)) &:= \langle V(X, A) \rangle_{p(A|X)} - \frac{1}{\beta} D_{\text{KL}}(p(A|X) || p_0(A)). \end{aligned}$$

Note that, here the expectation with respect to $p(X)$ does not affect the optimization with respect to $p(A|X)$ since it can be performed pointwise for each particular realization x of X . In fact, we would have obtained the same result when conditioning on an arbitrary value $X = x$

from the outset. However, in general, optimal information processing strategies may depend on the entire distribution $p(X)$ and can therefore not be obtained from only considering single observations x , for example when also optimizing with respect to the prior $p_0(A)$, see e.g., [85].

Free energy maximization. The optimal action distribution $p^*(A|X)$ maximizing F_A is a Boltzmann distribution (8) with “energy” $V(X, A)$ and prior $p_0(A)$,

$$p^*(A|X) = \frac{1}{\mathcal{Z}(X)} p_0(A) e^{\beta V(X,A)}, \quad (21)$$

where $\mathcal{Z}(X) := \sum_a p_0(a) e^{\beta V(X,a)}$. Note that in order to evaluate the utility V , it is required to determine the Bayes’ posterior $p(S|X)$. This shows how in a utility-based approach, the need to perform Bayesian inference results directly from the assumption about which variables are observed and which are not.

4.3 Critical points

The main idea of free energy in the context of information processing with limited resources is that any computation can be thought of abstractly as a transformation from a distribution p_0 of prior knowledge to a posterior distribution p that encapsulates an advanced state of knowledge resulting from deliberation. The progress that is made through such a transformation is quantitatively captured by two measures: the expected utility $\langle U \rangle_p$ that quantifies the quality of p and $C(p, p_0)$ that measures the cost of uncertainty reduction from p_0 to p . Clearly, the critical point of this framework is the choice of the cost function C . In particular, we could ask whether there is some kind of universal cost function that is applicable to any perception-action process or whether there are only problem-specific instantiations. Of course, having a universal measure that allows applying the same concepts to extremely diverse systems is both a boon and a bane, because the practical insights it may provide for any concrete instance could be very limited. This is the root of a number of critical issues:

- i. **What is the cost C ?** An important restriction of all deliberation costs of the form $C(p, p_0)$ is that they only depend on the initial and final distributions and ignore the process of how to get from p_0 to p . When varying a single resource (e.g., processing time) we can use $C(p, p_0)$ as a process-independent *proxy* for the resource. However, if there are multiple resources involved (e.g., processing time, memory, and power consumption), a single cost cannot tell us how these resources are weighted optimally without making further process-dependent assumptions. In general, the theory makes no suggestions whatsoever about mechanical processes that could implement resource-optimal strategies, it only serves as a baseline for comparison. Finally, simply requiring the measure to be monotonic in the uncertainty reduction, does not uniquely determine the form of C , as there have been multiple proposals of uncertainty measures in the literature (see e.g., [86]), where relative entropy is just one possibility. However, relative entropy is distinguished from all other uncertainty measures in its additivity property, that for example allows to express optimal probabilistic updates from p_0 to p in terms of additions or subtractions of utilities, such as log-likelihoods for evidence accumulation in Bayesian inference.
- ii. **What is the utility?** When systems are engineered, utilities are usually assumed to be given such that desired behavior is specified by utility maximization. However, when we observe perception-action systems, it is often not so clear what the utility should be, or in fact, whether there even exists a utility that captures the observed behavior in terms of utility maximization. This question of the identifiability of a utility function is studied extensively in the economic sciences, where the basic idea is that systems reveal their preferences

through their actual choices and that these preferences have to satisfy certain consistency axioms in order to guarantee the existence of a utility function. In practice, to guarantee unique identifiability these axioms are usually rather strong, for example ignoring the effects of history and context when choosing between different items, or ignoring the possibility that there might be multiple objectives. When not making these strong assumptions, utility becomes a rather generic concept, like the concept of probability, and additional assumptions like soft-maximization are necessary to translate from utilities to choice probabilities.

- iii. **The problem of infinite regress.** One of the main conceptual issues with the interpretation of C as a deliberation cost is that the original utility optimization problem is simply replaced by another optimization problem that may even be more difficult to solve. This novel optimization problem might again require resources to be solved and could therefore be described by a higher-level deliberation cost, thus leading to an infinite regress. In fact, any decision-making model that assumes that decision-makers reason about processing resources are affected by this problem [87, 88]. A possible way out is to consider the utility-information trade-off simply an *as if* description, since perception-action systems that are subject to a utility-information trade-off do not necessarily have to reason or know about their deliberation costs. It is straightforward, for example, to design processes that probabilistically optimize a given utility with no explicit notion of free energy, but for an outside observer the resulting choice distribution looks like an optimal free energy trade-off [89].

In summary, the free energy trade-off between utility and information primarily serves as a normative model for optimal probability assignments in information-processing nodes or networks. Like other Bayesian approaches, it can also serve as a guide for constructing and interpreting systems, although it is in general not a mechanistic model of behavior. In that respect it shares the fate of its cousins in thermodynamics and coding theory [90] in that they provide theoretical bounds on optimality but devise no mechanism for processes to achieve these bounds.

5 Variational free energy in Active Inference

5.1 The basic idea

Variational free energy is the main ingredient used in the *Free Energy Principle* for biological systems in the neuroscience literature [26, 33, 35, 91], which has been considered as “arguably the most ambitious theory of the brain available today” [92]. Since variational free energy in itself is just a mathematical construct to measure the dissimilarity between distributions and functions—see Section 3—, the biological content of the Free Energy Principle must come from somewhere else. The basic biological phenomenon that the Free Energy Principle purports to explain is *homeostasis*, the ability to actively maintain certain relevant variables (e.g., blood sugar) within a preferred range. Usually, homeostasis is applied as an explanatory principle in physiology whereby the actual value of a variable is compared to a target value and corrections to deviation errors are made through a feedback loop. However, homeostasis has also been proposed as an explanatory principle for complex behavior in the cybernetic literature [93–96]—for example, maintaining blood sugar may entail complex feedback loops of learning to hunt, to trade and to buy food. Crucially, being able to exploit the environment in order to attain favorable sensory states, requires implicit or explicit knowledge of the environment that could either be pre-programmed (e.g., insect locomotion) or learnt (e.g., playing the piano).

The Free Energy Principle was originally suggested as a theory of cortical responses [33] by promoting the free energy formulation of predictive coding that was introduced by Dayan and

Hinton with the Helmholtz machine [5]. It found its most recent incarnation in what is known as *Active Inference* that attempts to extend variational Bayesian inference to the problem of action selection. Here, the target value of homeostasis is expressed through a probability distribution p_{des} under which desired sensory states have a high probability. The required knowledge about the environment is expressed through a generative model p_0 that relates observations, hidden causes, and actions. As the generative model allows to make predictions about future states and observations, it enables to choose actions in such a way that the predicted consequences conform to the desired distribution. In Active Inference, this is achieved by merging the generative and the desired distributions, p_0 and p_{des} , into a single reference function ϕ to which trial distributions q over the unknown variables are fitted by minimizing the variational free energy $F(q||\phi)$. This free energy minimization is analogous to variational Bayesian inference, where the reference is always given by a joint distribution evaluated at observed quantities (cf. Section 3.2.1). In the resulting homeostatic process, the trial distributions q play the role of internal variables that are manipulated in order to achieve desired sensory consequences that are not directly controllable. Minimizing variational free energy by the alternating variation of trial distributions over actions $q_{Actions}$ and trial distributions over hidden states q_{States} ,

$$\underbrace{\min_{q_{Actions}} F(q||\phi)}_{\text{Action}} \quad \text{and} \quad \underbrace{\min_{q_{States}} F(q||\phi)}_{\text{Perception}}, \tag{22}$$

is then equated with processes of action and perception.

In a nutshell, the central tenet of the Free Energy Principle states that organisms maintain homeostasis through minimization of variational free energy between a trial distribution q and a reference function ϕ by acting and perceiving. Sometimes the even stronger statement is made that minimizing variational free energy is *mandatory* for homeostatic systems [97, 98].

5.2 A simple example

Ingredients. Applying the Active Inference recipe (cf. Fig 7) to our running example from Fig 1 with current and future states S, S' , current and future observations X, X' , and action A , we need a generative model p_0 , a desired distribution p_{des} , and trial distributions q . The generative model $p_0(\mathbf{X}, \mathbf{S}, A)$ is specified by the factors in the decomposition (1), the desired distribution $p_{des}(X')$ is a given fixed probability distribution over future sensory states X' , and the trial distributions q are probability distributions over all unknown variables, S, S', X' , and A .

In most treatments of Active Inference in the literature, the trial distributions q are simplified, either by a full mean-field approximation over states and actions [34, 35], by a partial mean-field approximation where the dependency on actions is kept but the states are treated independently of each other [99, 100], or more recently [101, 102] by the so-called Bethe approximation [58, 65], where subsequent states are allowed to interact. In the partial mean-field assumption of [99], the trial distribution over X' is fixed and given by $p_0(X'|S')$, while for A, S and S' the trial distributions are variable but restricted to be of the mean-field form for S and S' ,

$$q(\mathbf{S}, A) = q(S) q(S'|A) q(A), \tag{23}$$

i.e., the hidden states S and S' are assumed to be independent given A . While mean-field approximations can be good enough for simple perceptual inference, where a *single* hidden cause might be responsible for a set of observations, they can be too strong simplifications for sequential decision-making problems where the next state S' depends on the previous state S .

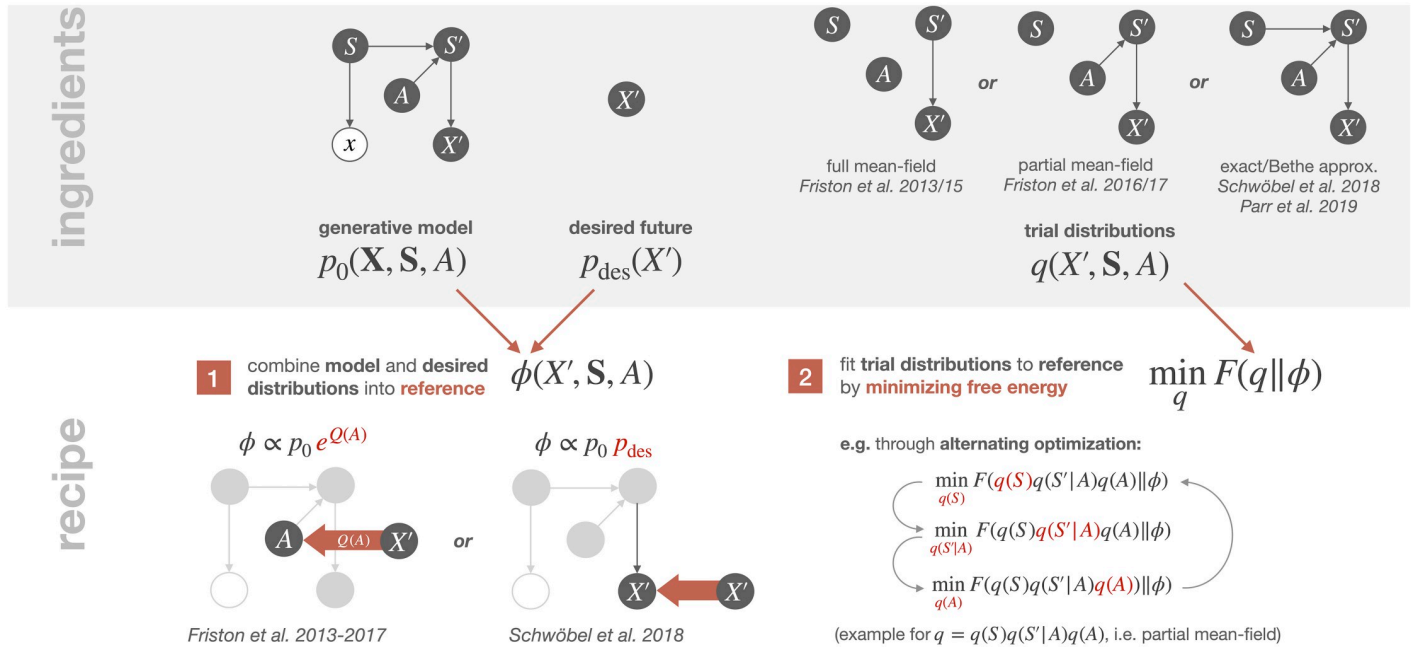


Fig 7. Overview of the Active Inference recipe, applied to our example from Fig 1.

<https://doi.org/10.1371/journal.pcbi.1008420.g007>

In fact, as can be seen for example in [S2 Notebook](#), mean-field assumptions may fail to show goal-directed behavior even for very simple tasks such as the navigation in a grid world. A less restrictive assumption would be a Bethe approximation, a special case of Kikuchi’s cluster variation method [103], which allows S and S' as well as S' and X' to be stochastically dependent—cf. Section C in [S1 Appendix](#), where we derive the update equations under the Bethe assumption for the simple example of this section. In general, the Bethe approximation achieves exact marginals in tree-like models, such as the models that are considered in the Active Inference literature, because it results in update equations that are equivalent to Pearl’s belief propagation algorithm [40, 58].

Reference function. The reference ϕ is constructed by combining the two distributions p_{des} and p_0 . To do so, there have been several proposals in the Active Inference literature, which fall into one of two categories: Either a specific value function Q is defined (containing p_{des}), which is multiplied to the generative model using a soft-max function [35, 99, 100],

$$\phi(X', \mathbf{S}, A) := p_0(X = x, X', \mathbf{S}|A) \frac{1}{Z} p_0(A) e^{Q(A)}, \tag{24}$$

or the desired distribution is multiplied directly to the generative model [101],

$$\phi(X', \mathbf{S}, A) := p_{des}(X') p_0(X = x, X', \mathbf{S}, A) \tag{25}$$

While the reference function in (25) is already completely specified, we still need to know how to determine the value function Q in the case of (24). For the partial mean-field assumption (23) it is defined in the literature [99, 100] as

$$Q(a) := \langle U(X', S') \rangle_{q(X', S'|A=a)} + H(q(X'|A = a)), \tag{26}$$

where $U(x', s') := \log p_{des}(x') + \log p_0(x'|s')$ favors both desirable and plausible future

observations x' . While here desirability and plausibility is built into the value function Q idiosyncratically, in utility-based approaches (cf. Section 4.2) only desirability has to be put into the design of the utility function, because there the likelihood $p_0(X'|S')$ of future observations is automatically taken into account by the expected utility V that is (soft-)maximized by (21). Moreover, since Q can be rewritten as

$$Q(a) = -D_{\text{KL}}(q(X'|A) \| p_{\text{des}}(X')) - \langle H(p_0(X'|S')) \rangle_{q(S'|A)},$$

the extra entropy term in (26) has the effect of actions leading to consequences that more or less *match* the desired distribution, while also explicitly punishing actions that lead to a high variability of observations (by requiring a low average entropy of $p_0(X'|S')$), rather than trying to produce the single most desired outcome—see the discussion at the end of Section 5.3. Note also that the value function Q depends (non-linearly) on the trial distribution $q(S'|A)$, because $q(X'|A) = \sum_{s'} p_0(X'|s')q(s'|A)$ is itself a function of $q(S'|A)$, which is problematic during free energy minimization (see (ii) in Section 5.3).

Free energy minimization. Once the form of the trial distributions q —e.g., by a partial mean-field assumption (23) or a Bethe approximation (see S1 Appendix)—and the reference ϕ are defined, the variational free energy is simply determined by $F(q|\phi)$. In the case of a mean-field assumption, the resulting free energy minimization problem is solved approximately by performing an alternating optimization scheme, in which the variational free energy is minimized separately with respect to each of the variable factors in a factorization of q , for example by alternating between $\min_{q(S)} F$, $\min_{q(S'|A)} F$, and $\min_{q(A)} F$ in the case of the partial mean-field assumption (23), where in each step the factors that are not optimized are kept fixed (cf. Fig 7). In S1 Appendix we derive the update equations for the cases (24) and (25) under mean-field and Bethe approximations for the one-step example discussed in this section. Mean-field solutions for the general case of arbitrarily many timesteps together with their exact solutions can be found in S1 Notebook, where we also highlight the theoretical differences between various proposed formulations of Active Inference. The effect of some of these differences can be seen in the grid world simulations in S2 Notebook.

5.3 Critical points

The main idea behind Active Inference is to express the problem of action selection in a similar manner to the perceptual problem of Bayesian inference over hidden causes. In Bayesian inference, agents are equipped with likelihood models $p_0(X|Z)$ that determine the desirability of different hypotheses Z under known data X . In Active Inference, agents are equipped with a given desired distribution $p_{\text{des}}(X')$ over future outcomes that ultimately determines the desirability of actions A . An important difference that arises is that perceptual inference has to condition on past observations $X = x$, whereas naive inference over actions would have to condition on desired future outcomes $X' = x'$.

For a single desired future observation x' , Bayesian inference could be applied in a straightforward way by simply conditioning the generative model p_0 on $X' = x'$. Similarly, one could condition on a desired distribution $p_{\text{des}}(X')$ using Jeffrey's conditioning rule [104], resulting in $p(A|p_{\text{des}}) = \sum_{x'} p(A|x') p_{\text{des}}(x')$, which could be implemented by first sampling a goal $x' \sim p_{\text{des}}(X')$ and then inferring $p(A|x')$ given the single desired observation x' . However, one of the problems with such a naive approach is that the choice of a goal is solely determined by its desirability, whereas its realizability for the decision-maker is not taken into account. This is because by conditioning on p_{des} , the decision-maker effectively seeks to choose actions in order to *reproduce* or *match* the desired distribution.

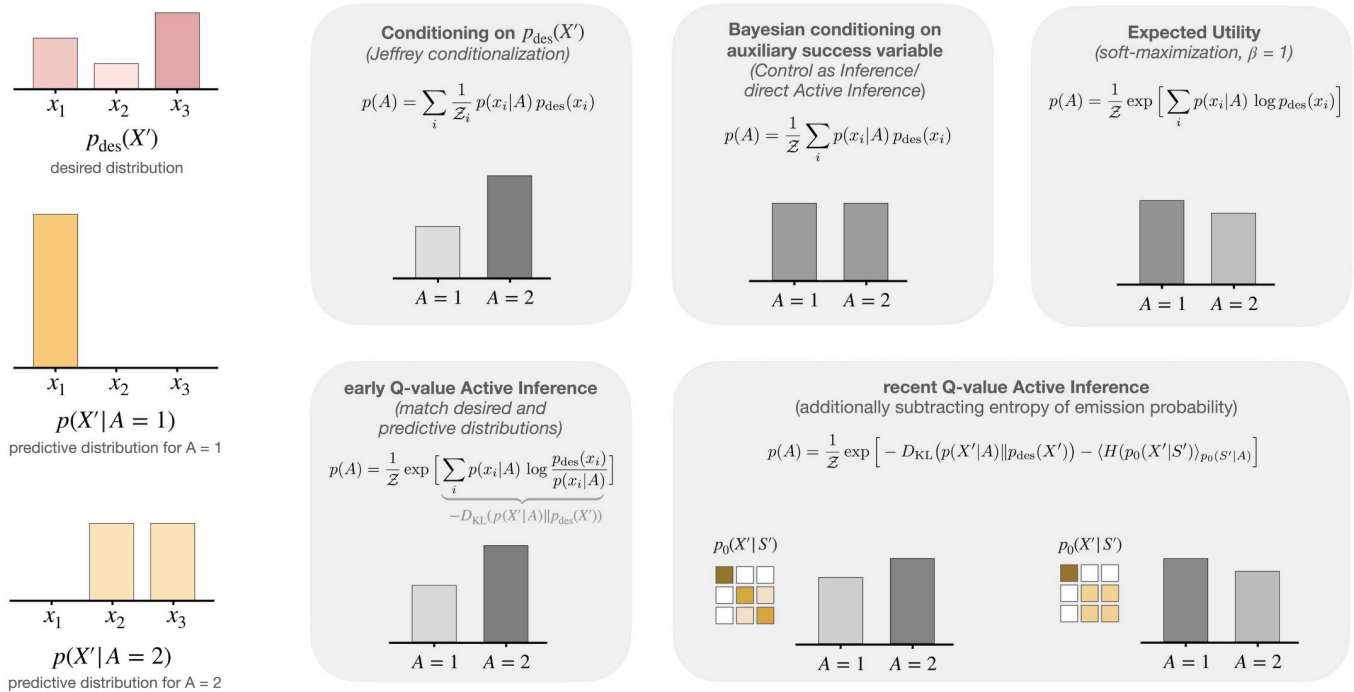


Fig 8. Consequences of assuming a desired distribution p_{des} for action planning under purely inference-based methods, expected utility, and Active Inference, in the case of a simple example with two actions, one with a deterministic outcome and one with random outcomes. As can be seen from the displayed equations, conditioning on p_{des} (Jeffrey conditionalization) and conditioning on success (Control as Inference/direct Active Inference) only differ in the order of normalizing and taking the expectation over X' . While conditioning on p_{des} requires to first sample a target outcome from p_{des} before an action from $p(A|X')$ can be planned, conditioning on success directly weighs the desirability of an outcome $p_{des}(x')$ by its realizability $p(x'|A)$. From this point of view, the expected utility approach is very similar to Control as Inference (which can also be seen in the grid world environment [S2 Notebook](#)), since it also weighs the utility of an outcome with its realizability before soft-maximizing. It only differs in how it treats the desired distribution as an exponentiated utility, moving the utility values closer together so that option $A = 1$ is slightly preferred. The early version [34] of Active Inference is similar to Jeffrey conditioning, because decision-makers are also assumed to *match* the desired distribution, by defining the value function Q as a KL divergence between the predicted and desired distributions. In later versions of Q-value Active Inference [35, 99, 100], the value function Q is modified by an additional entropy term that explicitly punishes observations with high variability. Consequently, even when the effect of the action on future observations is kept the same, i.e., the predictive distribution $p(X'|A) = \sum_{s'} p_0(X'|s')p_0(s'|A)$ remains as depicted in the left-hand column, the preference over actions now changes completely depending on $p_0(X'|S')$ —whereas in the other approaches, only the predictive distribution $p(X'|A)$ and $p_{des}(X')$ influence planning. While there might be circumstances where this extra punishment of high outcome variability could be beneficial, it is questionable from a normative point of view why anything else other than the predicted outcome probability $p(X'|A)$ should be considered for planning. See [S2 Appendix](#) for details about the choices made in the example.

<https://doi.org/10.1371/journal.pcbi.1008420.g008>

To overcome this problem, *Control as Inference* or *Planning as Inference* approaches in the machine learning literature [77, 105–108] do not directly condition on desired future observations but on future *success* by introducing an auxiliary binary random variable R such that $R = 1$ encodes the occurrence of desired outcomes. The auxiliary variable R comes with a probability distribution $p_0(R|X', \dots)$ that determines how well the outcomes satisfy desirability criteria of the decision-maker, usually defined in terms of the reward or utility attached to certain outcomes—see the discussion in (iii) below. The extra variable gives the necessary flexibility to infer successful actions by simply conditioning on $R = 1$. The advantage of such an approach over direct Jeffrey conditionalization given a desired distribution over future observations can be seen in the grid world simulations in [S2 Notebook](#), especially the ability of choosing a desired outcome that is not only desirable but also achievable—see also [Fig 8](#).

Active Inference tries to overcome the same problem of reconciling realizability and desirability, but without explicitly introducing extra random variables and without explicitly conditioning on the future. Instead, the desired distribution is combined with the generative model to form a new reference function ϕ such that the posteriors q^* resulting from the minimization

of the free energy $F(q||\phi)$ contain a baked-in tendency to reach the desired future encoded by ϕ . This approach is the root of a number of critical issues with current formulations of Active Inference:

i. **How to incorporate the desired distribution into the reference?**

Instead of using Bayesian conditioning directly in order to condition the generative model p_0 on the desired future, in Active Inference it is required that the reference ϕ contains the desired distribution in a way such that actions sampled from the resulting posterior model are more likely if they lead to the desired future. As can be seen already for the one-step case in (24) and (25), the method of how to incorporate the desired distribution into the reference function is not unique and does not follow from first principles. There have been essentially two different proposals in the literature on Active Inference of how to combine the two distributions p_{des} and p_0 into ϕ (cf. Fig 7): Either a hand-crafted value function Q is designed that specifically modifies the action probability of the generative model, or the probability over futures X' under the generative model p_0 is modified by directly multiplying p_{des} to the likelihood $p_0(X'|S')$. We discuss both of these proposals in (ii) and (iii) below.

ii. **Proposal 1: Q-value Active Inference [34, 35, 99, 100]**

In the most popular formulation of Active Inference, the probability over actions in the reference ϕ is defined by $\frac{1}{Z} p_0(A) e^{Q(A)}$, where the value function Q (also called the “expected free energy”) depends non-linearly on the trial distributions q , as can be seen exemplarily in (26) for the one-step case under the partial mean-field assumption of [99, 100], where $q(S'|A)$ enters Q through $q(X'|A) = \sum_{s'} p_0(X'|s') q(s'|A)$. Note that, because of this non-linearity the alternating free energy minimization would have no closed-form solutions (cf. S1 Appendix). This means that both the trial distributions q and the reference $\phi = \phi(q)$ will change when q is varied during the minimization of the *total* variational free energy $F(q||\phi(q))$, as would be required when stipulating a single free energy functional for optimization. This highlights an important conceptual difference to variational Bayesian inference, where one assumes a *fixed* reference ϕ —resulting from the evaluation of a *fixed* probabilistic model p_0 at known variables (see Section 3.2.1)—to which distributions q are fitted by minimizing $F(q||\phi)$. In contrast, when changing the reference $\phi(q)$ during the optimization process, it is no longer clear what is actually achieved by this minimization. As demonstrated by S2 Notebook, this issue has immediate practical implications, as respecting or ignoring the extra q dependency can result in very different behavior even in simple grid world simulations.

In the Active Inference literature, however, the extra q -dependency of Q is largely ignored. Instead of optimizing the full free energy $F(q||\phi(q))$ with respect to state and action distributions, one alternately optimizes the free energy over states F_A for each action A and then the full free energy with respect to action distributions only, so that action and perception effectively optimize two different free energies. It is crucial to note, however, that unlike in variational Bayesian inference with fixed reference, this separation does not follow from the formalism of variational free energy, but is a design choice of the Active Inference framework that imposes this separation by force (see S4 Appendix for more details). This way, both separate optimizations can be considered as variational inference in each single update, even though when alternating them the reference ϕ still changes across the combined optimization process. This is in contrast to alternating optimization schemes in variational inference (e.g., in the Bayesian EM algorithm) where the reference ϕ does not change between optimization steps. Thus, there are two choices: Either Q-value Active Inference is regarded as some kind of approximation to variational inference under a single total free

energy, or one has to give up the idea of a single free energy function that is optimized. Either way, the combined process of action and perception does not correspond to a single variational inference process.

Finally, another important practical issue with Q-value Active Inference models is that the definition of Q relies on a mean-field approximation of the trial distributions q , under which hidden states are assumed to be stochastically independent. This simplification is too strong for sequential decision-making tasks, which renders the approach unfit for environments where the current state depends stochastically on previous states (see [S2 Notebook](#) for a demonstration).

iii. **Proposal 2: direct Active Inference** [101]

When multiplying p_{des} to the generative model directly, as in (25), then the resulting reference ϕ is no longer given by a joint distribution of observations, states, and actions (since in general $\sum_{x'} p_{\text{des}}(x')p_0(x'|S') \neq 1$). Instead, this formulation of Active Inference turns out to be a special case of previous Control as Inference approaches in the machine learning literature [105, 107], where one conditions on an auxiliary success variable R . In particular, for our running example from Fig 1 with a probabilistic model of the form (1), Control as Inference defines

$$p_0(R = 1|X', S', A) := e^{r(X', S', A)} = 1 - p_0(R = 0|X', S', A),$$

where $r = r(X', S', A)$ denotes a general (negative) reward function determining desirability. The full joint of the new set of variables is then given by

$$p_0(R, \mathbf{X}, \mathbf{S}, A) = p_0(R|X', S', A) p_0(\mathbf{X}, \mathbf{S}, A). \tag{27}$$

Control as Inference then conditions actions on both, the history and future success ($R = 1$). For our one-step example, this results in the Bayes' posterior

$$p(A|X = x, R = 1) = \frac{1}{Z} \sum_{x', s, s'} p_0(R = 1|x', s', A) p_0(\mathbf{x}, \mathbf{s}, A). \tag{28}$$

It is straightforward to identify $p_{\text{des}}(X')$ of Active Inference as a particular choice of a success probability $p_0(R = 1|X')$, or equivalently, $\log p_{\text{des}}(X')$ as a reward function $r = r(X')$, so that the joint distribution (27) reduces to the reference function ϕ in (25). Thus, the version of Active Inference in [101] is simply a variational formulation of Control as Inference that approximates exact posteriors of the form (28), like other previous variational Bayes' approaches [107, 109, 110].

In summary, the assumption of a desired distribution p_{des} over future outcomes has led to various attempts in the Active Inference literature of using probabilistic inference to determine profitable actions. Either an action distribution $\frac{1}{Z} p_0(A)e^{Q(A)}$ is built into the reference function, which presupposes optimal behavior by designing a value function Q that leads to desired consequences, or the outcome probability under the generative model p_0 is modified directly by multiplying p_{des} to p_0 . The latter case is the variational version of Control as Inference, well-known in the machine learning literature [77, 105–110]. Considering the issues of Q-value Active Inference discussed above, and the fact that Control as Inference does not rely on a desired distribution over outcomes, we could ask whether formulating preferences by assuming a desired distribution is well-advised. As can be seen from Fig 8, the difference between purely inference-based methods, expected utility approaches, and Active Inference is mainly in how they treat the desired distribution. Should p_{des} be matched or is it good enough if

actions are chosen that lead to a high desired outcome probability? While Control as Inference and utility-based models essentially take the latter approach, Q-value Active Inference answers this question by requiring that the desired distribution should be matched as long as the average entropy of $p_0(X|S)$ is small.

6 So what does free energy bring to the table?

6.1 A practical tool

It is unquestionable that the concept of free energy has seen many fruitful practical applications outside of physics in the statistical and machine learning literature. As has been discussed in Section 3, these applications generally fall into one of two categories, the principle of maximum entropy, and a variational formulation of Bayesian inference. Here, the principle of maximum entropy is interpreted in a wider sense of optimizing a trade-off between uncertainty (entropy) and the expected value of some quantity of interest (energy), which in practice often appears in the form of regularized optimization problems (e.g., to prevent overfitting) or as a general inference method allowing to determine unbiased priors and posteriors (cf. Section 3.1). In the variational formulation of Bayes' rule, free energy plays the role of an error measure that allows to do approximate inference by constraining the space of distributions over which free energy is optimized, but can also inform the design of efficient iterative inference algorithms that result from an alternating optimization scheme where in each step the full variational free energy is optimized only partially, such as the Bayesian EM algorithm, belief propagation, and other message passing algorithms (cf. Section 3.2).

It is important to realize that, while the mathematical expressions of a free energy from constraints with “energy” \mathcal{E} and trade-off parameter β and a variational free energy with reference ϕ can formally be transformed into each other by $\phi = e^{-\beta\mathcal{E}}$, the two kinds of free energy are inherently distinct, both methodically and by their motivation. In the case of the free energy from constraints, we are given a constraint on some quantity \mathcal{E} and we are trying to fulfil this constraint with minimum bias by selecting a distribution that trades off the two competing terms \mathcal{E} and entropy. This trade-off also gives the reason for the existence of the Lagrange multiplier β that has to be determined according to the constraint. In this sense the free energy from constraints is just a special case of the far more general Lagrangian method when applied to the optimization of expected values $\langle \mathcal{E} \rangle_p$ under entropy constraints (or the other way around). In contrast, variational free energy is simply a tool to represent the normalization of a reference function ϕ in terms of an optimization problem, and therefore does a priori not assume the existence of some quantity \mathcal{E} that we may have observed in an experiment or that has any other constraints attached, nor does one explicitly consider entropy to be constrained or optimized. Therefore, even though starting from a (positive) reference function ϕ we can always invent the existence of some quantity \mathcal{E} and some multiplier β such that $\phi = e^{-\beta\mathcal{E}}$, this does not explain why these quantities should exist or why they should be mapped into each other in that particular way. The Lagrangian method, on the other hand, explains why for a given constraint on \mathcal{E} we have a Lagrange multiplier β , how it is determined, and why the equilibrium distribution has the form $p^* = \frac{1}{Z} e^{-\beta\mathcal{E}}$.

6.2 Theories of intelligent agency

These practical use-cases of free energy formulations have also influenced models of intelligent behavior. In the cognitive and behavioral sciences, intelligent agency has been modelled in a number of different frameworks, including logic-based symbolic models, connectionist models, statistical decision-making models, and dynamical systems approaches. Even though

statistical thinking in a broader sense can in principle be applied to any of the other frameworks as well, statistical models of cognition in a more narrow sense have often focused on Bayesian models, where agents are equipped with probabilistic models of their environment allowing them to infer unknown variables in order to select actions that lead to desirable consequences [14, 76, 111]. Naturally, the inference of unknown variables in such models can be achieved by a plethora of methods including the two types of free energy approaches of maximum entropy and variational Bayes. However, both free energy formulations go one step further in that they attempt to extend both principles from the case of inference to the case of action selection: utility optimization with information constraints based on free energy from constraints and Active Inference based on variational free energy.

While sharing similar mathematical concepts, both approaches differ in syntax and semantics. An apparent apple of discord is the concept of utility [112]. Utility optimization with information constraints requires the determination of a utility function, whereas Active Inference requires the determination of a reference function. In the economic literature, subjective utility functions that quantify the preferences of decision-makers are typically restrictive in order to ensure identifiability when certain consistency axioms are satisfied. In contrast, in Active Inference the reference function involves determining a desired distribution given by the preferred frequency of outcomes. However, these differences start to vanish when weakening the utility concept to something like log-probabilities, such that the utility framework becomes more similar to the concept of probability that is able to explain arbitrary behavior. Moreover, Active Inference has to solve the additional problem of marrying up the agent's probabilistic model with its desired distribution into a single reference function (cf. Section 5.3). The solution to this problem is not unique, in particular it lies outside the scope of variational Bayesian inference, but it is critical for the resulting behavior because it determines the exact solutions that are approximated by free energy minimization. In fact, as can be seen in simple simulations such as [S2 Notebook](#), the various proposals for this merging that can be found in the Active Inference literature behave very differently.

Also, both approaches differ fundamentally in their motivation. The motivation of utility optimization with information constraints is to capture the trade-off between precision and uncertainty that underlies information processing. This trade-off takes the form of a free energy once an informational cost function has been chosen (cf. Section 4.3). Note that Bayes' rule can be seen as the minimum of a free energy from constraints with log-likelihoods as utilities, even though this equivalence is not the primary motivation of this trade-off. In contrast, Active Inference is motivated from casting the problem of action selection itself as an inference process [34], as this allows to express both action and perception as the result of minimizing the same function, the variational free energy. However, there is no mystery in having such a single optimization function, because the underlying probabilistic model already contains both action and perception variables in a single functional format and the variational free energy is just a function of that model. Moreover, while approximate inference can be formulated on the basis of variational free energy, inference in general does not rely on this concept, in particular inference over actions can easily be done without free energy [77, 105–107, 113].

However, there are also plenty of similarities between the two free energy approaches. For example, the assumption of a soft-max action distribution in Active Inference is similar to the posterior solutions resulting from utility optimization with information constraints. Moreover, the assumption of a desired future *distribution* relates to constrained computational resources, because the uncertainty constraint in a desired distribution over future states may not only be a consequence of environmental uncertainty, but could also originate from stochastic preferences of a satisficing decision-maker that accepts a wide range of outcomes. In fact, as we have seen in the discussion around [Fig 8](#), various methods for inference over actions differ in how

they treat preferences given by a distribution over desired outcomes: Some of them try to match the predictive and desired distributions, while others simply seek to reach states whose outcomes have a high desired probability. In [S2 Notebook](#), we provide a comparison of the discussed methods using grid world simulations, in order to see their resulting behavior also in a sequential decision-making task.

A remarkable resemblance among both approaches is the exclusive appearance of relative entropy to measure dissimilarity. In the Active Inference literature it is often claimed that every homeostatic system must minimize variational free energy [97], which is simply an extension of relative entropy for non-normalized reference functions (cf. Section 3.2.2). In utility-based approaches, the relative entropy (19) is typically used to measure the amount of information processing, even though theoretically other cost functions would be conceivable [74]. For a given homeostatic process, the KL divergence measures the dissimilarity between the current distribution and the limiting distribution and therefore is reduced while approximating the equilibrium. Similarly, in utility-based decision-making models, relative entropy measures the dissimilarity between the current posterior and the prior. In the Active Inference literature the stepwise minimization of variational free energy that goes along with KL minimization is often equated with the minimization of sensory *surprise* (see [S3 Appendix](#) for a more detailed explanation), an idea that stems from maximum likelihood algorithms, but that has been challenged as a general principle (see [114] and the response [115]). Similarly, one could in principle rewrite free energy from constraints in terms of informational surprise, which would however simply be a rewording of the probabilistic concepts in log-space. The same kind of rewording is well-known between probabilistic inference and the minimum description length principle [116] that also operates in log-space, and thus reformulates the inference problem as a surprise minimization problem without adding any new features or properties.

6.3 Biological relevance

So far we have seen how free energy is used as a technical instrument to solve inference problems and its corresponding appearance in different models of intelligent agency. Crucially, these kinds of models can be applied to any input-output system, be it a human that reacts to sensory stimuli, a cell that tries to maintain homeostasis, or a particle trapped by a physical potential. Given the existing literature that has widely applied the concept of free energy to biological systems, we may ask whether there are any specific biological implications of these models.

Considering free energy from constraints, the trade-off between utility and information processing costs provides a normative model of decision-making under resource constraints, that extends previous optimality models based on expected utility maximization and Bayesian inference. Analogous to rate-distortion curves in information theory, optimal solutions to decision-making problems are obtained that separate achievable from non-achievable regions in the information-utility plane (cf. [Fig 5](#)). The behavior of real decision-making systems under varying information constraints can be analyzed experimentally by comparing their performance with respect to the corresponding optimality curve. One can experimentally relate abstract information processing costs measured in bits to task-dependent resource costs like reaction or planning times [20, 22]. Moreover, the free energy trade-off can also be used to describe networks of agents, where each agent is limited in its ability, but the system as a whole has a higher information processing capacity—for example, neurons in a brain or humans in a group. In such systems different levels of abstraction arise depending on the different positions of decision-makers in the network [23, 71, 85]. As we have discussed in Section 4.3, just like coding and rate-distortion theory, utility theory with information costs can only provide

optimality bounds but does not specify any particular mechanism of how to achieve optimality. However, by including more and more constraints one can make a model more and more mechanistic and thereby gradually move from a normative to a more descriptive model, such as models that consider the communication channel capacity of neurons with a finite energy budget [24].

Considering variational free energy, there is a vast literature on biological applications mostly focusing on neural processing (e.g., predictive coding and dopamine) [102, 117, 118], but there are also a number of applications aiming to explain behavior (e.g., human decision-making and hallucinations) [119]. Similarly to utility-based models, Active Inference models can be studied in terms of *as if* models, so that actual behavior can be compared to predicted behavior as long as suitable prior and likelihood models can be identified from the experiment. When applied to brain dynamics, the *as if* models are sometimes also given a mechanistic interpretation by relating iterative update equations that appear when minimizing variational free energy with dynamics in neuronal circuits. As discussed in Section 3.2.3, the update equations resulting for example from mean-field or Bethe approximations, can often be written in message passing form in the sense that the update for a given variable only has contributions that requires the current approximate posterior of neighbouring nodes in the probabilistic model. These contributions are interpreted as local messages passed between the nodes and might be related to brain signals [102]. Other interpretations [28, 91, 100] obtain similar update equations by minimizing variational free energy directly through gradient descent, which can again be related to neural coding schemes like predictive coding. As these coding schemes have existed irrespective of free energy [120, 121], especially since minimization of prediction errors is already seen in maximum likelihood estimation [120], the question remains whether there are any specific predictions of the Active Inference framework that cannot be explained with previous models (see [39, 122] for recent discussions of this question).

6.4 Conclusion

Any theory about intelligent behavior has to answer three questions: *Where am I?*, *where do I want to go?*, and *how do I get there?*, corresponding to the three problems of inference and perception, goals and preferences, and planning and execution. All three problems can be addressed either in the language of probabilities or utilities. Perceptual inference can either be considered as finding parameters that maximize probabilities or likelihood utilities. Goals and preferences can either be expressed by utilities over outcomes or by desired distributions. The third question can be answered by the two free energy approaches that either determine future utilities based on model predictions, or infer actions that lead to outcomes predicted to have high desired probability or match the desired distribution. In standard decision-making models actions are usually determined by a utility function that ranks different options, whereas perceptual inference is determined by a likelihood model that quantifies how probable certain observations are. In contrast, both free energy approaches have in common that they treat all types of information processing, from action planning to perception, as the same formal process of minimizing some form of free energy. But the crucial difference is not whether they use utilities or probabilities, but how predictions and goals are interwoven into action.

This article started out by tracing back the seemingly mysterious connection between Helmholtz free energy from thermodynamics and Helmholtz' view of model-based information processing that led to the analysis-by-synthesis approach of perception, as exemplified in predictive coding schemes, and in particular to discuss the role of free energy in current models of intelligent behavior. The mystery starts to dissolve when we consider the two kinds of free energies discussed in this article, one based on the maximum entropy principle and the

other based on variational free energy—a dissimilarity measure between distributions and (generally unnormalized) functions that extends the well-known KL divergence from information theory. The Helmholtz free energy is a particular example of an energy information trade-off that results from the maximum entropy principle [46]. Analysis-by-synthesis is a particular application of inference to perception, where determining model parameters and hidden states can either be seen as a result of maximum entropy under observational constraints or of fitting parameter distributions to the model through variational free energy minimization. Thus, both notions of free energy can be formally related as entropy-regularized maximization of log-probabilities.

Conceptually, however, utility-based models with information constraints serve primarily as *ultimate* explanations of behavior, this means they do not focus on mechanism, but on the goals of behavior and their realizability under ideal circumstances. They have the appeal of being relatively straightforward generalization of standard utility theory, but they rely on abstract concepts like utility and relative entropy that may not be so straightforwardly related to experimental settings. While these normative models have no immediate mechanistic interpretation, their relevance for mechanistic models may be analogous to the relevance of optimality bounds in Shannon's information theory for practical codes [90]. In contrast, Active Inference models of behavior often mix ultimate and *proximate* arguments of explaining behavior [123, 124], because they combine the normative aspect of optimizing variational free energy with the mechanistic interpretation of the particular form of approximate solutions to this optimization. While mean-field approaches of Active Inference may be particularly amenable to such mechanistic interpretations, they are often too simple to capture complex behavior. In contrast, the solutions of direct Active Inference resulting from a Bethe assumption are equivalent to previous Control as Inference approaches [77, 105–110] that allow for Bayesian message passing formulations whose biological implementability can be debated irrespective of the existence of a free energy functional.

Finally, both kinds of free energy formulations of intelligent agency are so general and flexible in their ingredients that it might be more appropriate to consider them languages or tools to phrase and describe behavior rather than theories that explain behavior, in a sense similar to how statistics and probability theory are not biological or physical theories but simply provide a language in which we can phrase our biological and physical assumptions.

Supporting information

S1 Appendix. Derivation of exemplary update equations. We derive update equations of Q -value and direct Active Inference for the example in Section 5.2 under mean-field and Bethe approximations.

(PDF)

S2 Appendix. Uncertain and deterministic options. We give additional details on the example shown in Fig 8 that illustrates the effects of assuming a particular desired distribution over three outcomes under Jeffrey conditionalization, Control as Inference, expected utility optimization, and Active Inference.

(PDF)

S3 Appendix. Surprise minimization. Explanation of the relation between free energy minimization, free energy as a bound on surprise, and surprise minimization.

(PDF)

S4 Appendix. Separation of model and state variables. Discussion of how model and state variables can be separated in variational Bayesian inference which motivates the optimization

scheme chosen by Active Inference.
(PDF)

S1 Notebook. Comparison of different formulations of Active Inference. A detailed comparison of the different formulations of Active Inference found in the literature (2013–2019), including their mean-field and exact solutions in the general case of arbitrary many time steps.
(HTML)

S2 Notebook. Grid world simulations. We provide implementations of the models discussed in this article in a grid world environment, both as a rendered html file as well as a jupyter notebook that is [available on github](#).
(HTML)

References

1. Yuille A, Kersten D. Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*. 2006; 10(7):301–308. <https://doi.org/10.1016/j.tics.2006.05.002> PMID: 16784882
2. Kawato M. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*. 1999; 9(6):718–727. [https://doi.org/10.1016/S0959-4388\(99\)00028-8](https://doi.org/10.1016/S0959-4388(99)00028-8) PMID: 10607637
3. Flanagan JR, Vetter P, Johansson RS, Wolpert DM. Prediction Precedes Control in Motor Learning. *Current Biology*. 2003; 13(2):146–150. [https://doi.org/10.1016/S0960-9822\(03\)00007-1](https://doi.org/10.1016/S0960-9822(03)00007-1) PMID: 12546789
4. Doya K. *Bayesian Brain: Probabilistic Approaches to Neural Coding*. Cambridge, Mass: MIT Press; 2007.
5. Dayan P, Hinton GE, Neal RM, Zemel RS. The Helmholtz Machine. *Neural Comput*. 1995; 7(5):889–904. <https://doi.org/10.1162/neco.1995.7.5.889> PMID: 7584891
6. Neal RM, Hinton GE. A View of the EM Algorithm that Justifies Incremental, Sparse, and other Variants. In: Jordan MI, editor. *Learning in Graphical Models*. Dordrecht: Springer Netherlands; 1998. p. 355–368.
7. Beal MJ. *Variational Algorithms for Approximate Bayesian Inference*. University of Cambridge, UK; 2003.
8. Williams RJ, Peng J. Function Optimization using Connectionist Reinforcement Learning Algorithms. *Connection Science*. 1991; 3(3):241–268. <https://doi.org/10.1080/09540099108946587>
9. Mnih V, Badia AP, Mirza M, Graves A, Lillicrap T, Harley T, et al. Asynchronous Methods for Deep Reinforcement Learning. In: Balcan MF, Weinberger KQ, editors. *Proceedings of The 33rd International Conference on Machine Learning*. vol. 48 of *Proceedings of Machine Learning Research*. New York, New York, USA: PMLR; 2016. p. 1928–1937. <http://proceedings.mlr.press/v48/mniha16.html>.
10. McKelvey RD, Palfrey TR. Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior*. 1995; 10(1):6–38. <https://doi.org/10.1006/game.1995.1023>
11. Sims CA. Implications of rational inattention. *Journal of Monetary Economics*. 2003; 50(3):665–690. [https://doi.org/10.1016/S0304-3932\(03\)00029-1](https://doi.org/10.1016/S0304-3932(03)00029-1)
12. Mattsson LG, Weibull JW. Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior*. 2002; 41(1):61–78. [https://doi.org/10.1016/S0899-8256\(02\)00014-3](https://doi.org/10.1016/S0899-8256(02)00014-3)
13. McFadden DL. Revealed stochastic preference: a synthesis. *Economic Theory*. 2005; 26(2):245–264. <https://doi.org/10.1007/s00199-004-0495-3>
14. Wolpert DH. In: *Information Theory—The Bridge Connecting Bounded Rational Game Theory and Statistical Physics*. Springer Berlin Heidelberg; 2006. p. 262–290.
15. Maccheroni F, Marinacci M, Rustichini A. Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica*. 2006; 74(6):1447–1498. <https://doi.org/10.1111/j.1468-0262.2006.00716.x>
16. Hansen LP, Sargent TJ. *Robustness*. Princeton University Press; 2008.
17. Still S. Information-theoretic approach to interactive learning. *Europhysics Letters*. 2009; 85(2):28005. <https://doi.org/10.1209/0295-5075/85/28005>
18. Tishby N, Polani D. Information Theory of Decisions and Actions. In: Cutsuridis V, Hussain A, Taylor JG, editors. *Perception-Action Cycle: Models, Architectures, and Hardware*. Springer New York; 2011. p. 601–636.

19. Ortega PA, Braun DA. Thermodynamics as a theory of decision-making with information-processing costs. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 2013; 469(2153):20120683. <https://doi.org/10.1098/rspa.2012.0683>
20. Ortega PA, Stocker A. Human Decision-Making under Limited Time. In: 30th Conference on Neural Information Processing Systems; 2016.
21. Sims CR. Rate–distortion theory and human perception. *Cognition*. 2016; 152:181–198. <https://doi.org/10.1016/j.cognition.2016.03.020> PMID: 27107330
22. Schach S, Gottwald S, Braun DA. Quantifying Motor Task Performance by Bounded Rational Decision Theory. *Frontiers in Neuroscience*. 2018; 12:932. <https://doi.org/10.3389/fnins.2018.00932> PMID: 30618561
23. Lindig-León C, Gottwald S, Braun DA. Analyzing Abstraction and Hierarchical Decision-Making in Absolute Identification by Information-Theoretic Bounded Rationality. *Frontiers in Neuroscience*. 2019; 13:1230. <https://doi.org/10.3389/fnins.2019.01230> PMID: 31824241
24. Bhui R, Gershman SJ. Decision by sampling implements efficient coding of psychoeconomic functions. *Psychological Review*. 2018; 125(6):985–1001. <https://doi.org/10.1037/rev0000123> PMID: 30431303
25. Ho MK, Abel D, Cohen JD, Littman ML, Griffiths TL. The Efficiency of Human Cognition Reflects Planned Information Processing. *Proceedings of the 34th AAAI Conference on Artificial Intelligence*. 2020;.
26. Friston KJ. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*. 2010; 11:127–138. <https://doi.org/10.1038/nrn2787> PMID: 20068583
27. Sales AC, Friston KJ, Jones MW, Pickering AE, Moran RJ. Locus Coeruleus tracking of prediction errors optimises cognitive flexibility: An Active Inference model. *PLOS Computational Biology*. 2019; 15(1):e1006267. <https://doi.org/10.1371/journal.pcbi.1006267> PMID: 30608922
28. Bogacz R. A tutorial on the free-energy framework for modelling perception and learning. *Journal of Mathematical Psychology*. 2017; 76:198–211. <https://doi.org/10.1016/j.jmp.2015.11.003> PMID: 28298703
29. Friston KJ, Shiner T, FitzGerald T, Galea JM, Adams R, Brown H, et al. Dopamine, Affordance and Active Inference. *PLoS Computational Biology*. 2012; 8(1):e1002327. <https://doi.org/10.1371/journal.pcbi.1002327> PMID: 22241972
30. Parr T, Friston KJ. Working memory, attention, and salience in active inference. *Scientific reports*. 2017; 7(1):14678–14678. <https://doi.org/10.1038/s41598-017-15249-0> PMID: 29116142
31. Mirza MB, Adams RA, Mathys C, Friston KJ. Human visual exploration reduces uncertainty about the sensed world. *PLOS ONE*. 2018; 13(1):e0190429. <https://doi.org/10.1371/journal.pone.0190429> PMID: 29304087
32. Parr T, Friston KJ. Generalised free energy and active inference. *Biological Cybernetics*. 2019; <https://doi.org/10.1007/s00422-019-00805-w> PMID: 31562544
33. Friston KJ. A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2005; 360(1456):815–836. <https://doi.org/10.1098/rstb.2005.1622> PMID: 15937014
34. Friston K, Schwartenbeck P, Fitzgerald T, Moutoussis M, Behrens T, Dolan R. The anatomy of choice: active inference and agency. *Frontiers in Human Neuroscience*. 2013; 7:598. <https://doi.org/10.3389/fnhum.2013.00598> PMID: 24093015
35. Friston KJ, Rigoli F, Ognibene D, Mathys C, Fitzgerald T, Pezzulo G. Active inference and epistemic value. *Cognitive Neuroscience*. 2015; 6(4):187–214. <https://doi.org/10.1080/17588928.2015.1020053> PMID: 25689102
36. Schwartenbeck P, Friston K. Computational Phenotyping in Psychiatry: A Worked Example. *eNeuro*. 2016; 3(4): <https://doi.org/10.1523/ENEURO.0049-16.2016> PMID: 27517087
37. Linson A, Parr T, Friston KJ. Active inference, stressors, and psychological trauma: A neuroethological model of (mal)adaptive explore-exploit dynamics in ecological context. *Behavioural Brain Research*. 2020; 380:112421. <https://doi.org/10.1016/j.bbr.2019.112421> PMID: 31830495
38. Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*. 2013; 36(3):181–204. <https://doi.org/10.1017/S0140525X12000477> PMID: 23663408
39. Colombo M, Wright C. First principles in the life sciences: the free-energy principle, organicism, and mechanism. *Synthese*. 2018;
40. Pearl J. Belief Updating by Network Propagation. In: Pearl J, editor. *Probabilistic Reasoning in Intelligent Systems*. San Francisco (CA): Morgan Kaufmann; 1988. p. 143–237.

41. Minka TP. Expectation Propagation for Approximate Bayesian Inference. In: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence. UAI'01. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.; 2001. p. 362–369.
42. Hinton GE, van Camp D. Keeping the Neural Networks Simple by Minimizing the Description Length of the Weights. In: Proceedings of the Sixth Annual Conference on Computational Learning Theory. COLT'93. New York, NY, USA: ACM; 1993. p. 5–13.
43. MacKay DJC. Information Theory, Inference & Learning Algorithms. USA: Cambridge University Press; 2002.
44. Boutilier C, Dean T, Hanks S. Decision-Theoretic Planning: Structural Assumptions and Computational Leverage. *J Artif Int Res.* 1999; 11(1):1–94.
45. Feynman RP, Hey AJG, Allen RW. Feynman Lectures on Computation. Advanced book program. Addison-Wesley; 1996.
46. Jaynes ET. Information Theory and Statistical Mechanics. *Phys Rev.* 1957; 106:620–630. <https://doi.org/10.1103/PhysRev.106.620>
47. Jaynes ET. Probability Theory. Bretthorst GL, editor. Cambridge University Press; 2003.
48. Rosenkrantz RD. E.T. Jaynes: Papers on Probability, Statistics and Statistical Physics. Dordrecht: Springer Netherlands; 1983.
49. Bernoulli J. *Ars conjectandi*. Basel, Thurneysen Brothers; 1713.
50. de Laplace PS. *Théorie analytique des probabilités*. Ve. Courcier, Paris; 1812.
51. Poincaré H. *Calcul des probabilités*. Gauthier-Villars, Paris; 1912.
52. Williams PM. Bayesian Conditionalisation and the Principle of Minimum Information. *The British Journal for the Philosophy of Science.* 1980; 31(2):131–144. <https://doi.org/10.1093/bjps/31.2.131>
53. Haarnoja T, Tang H, Abbeel P, Levine S. Reinforcement Learning with Deep Energy-Based Policies. In: ICML; 2017.
54. Fox R, Pakman A, Tishby N. Taming the Noise in Reinforcement Learning via Soft Updates. In: Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence. UAI'16. Arlington, Virginia, United States: AUA Press; 2016. p. 202–211. <http://dl.acm.org/citation.cfm?id=3020948.3020970>.
55. Koller D. Probabilistic graphical models: principles and techniques. Cambridge, Massachusetts: The MIT Press; 2009.
56. Opper M, Saad D. In: Comparing the Mean Field Method and Belief Propagation for Approximate Inference in MRFs; 2001. p. 229–239.
57. Dempster AP, Laird NM, Rubin DB. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society Series B (Methodological).* 1977; 39(1):1–38. <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x>
58. Yedidia JS, Freeman WT, Weiss Y. Generalized Belief Propagation. In: Leen TK, Dietterich TG, Tresp V, editors. *Advances in Neural Information Processing Systems 13*. MIT Press; 2001. p. 689–695.
59. Wainwright MJ, Jaakkola TS, Willsky AS. MAP estimation via agreement on (hyper)trees: Message-passing and linear-programming approaches. *IEEE Transactions on Information Theory.* 2005; 51(11):3697–3717. <https://doi.org/10.1109/TIT.2005.856938>
60. Winn J, Bishop CM. Variational Message Passing. *J Mach Learn Res.* 2005; 6:661–694.
61. Minka T. Divergence Measures and Message Passing. Microsoft; 2005. MSR-TR-2005-173.
62. Yedidia JS, Freeman WT, Weiss Y. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Transactions on Information Theory.* 2005; 51(7):2282–2312. <https://doi.org/10.1109/TIT.2005.850085>
63. Csiszár I, Tusnády G. Information geometry and alternating minimization procedures. *Statistics and Decisions, Supplement Issue.* 1984; 1:205–237.
64. Hathaway RJ. Another interpretation of the EM algorithm for mixture distributions. *Statistics & Probability Letters.* 1986; 4(2):53–56. [https://doi.org/10.1016/0167-7152\(86\)90016-7](https://doi.org/10.1016/0167-7152(86)90016-7)
65. Heskes T. Stable Fixed Points of Loopy Belief Propagation Are Local Minima of the Bethe Free Energy. In: Becker S, Thrun S, Obermayer K, editors. *Advances in Neural Information Processing Systems 15*. MIT Press; 2003. p. 359–366.
66. Yuille AL. CCCP Algorithms to Minimize the Bethe and Kikuchi Free Energies: Convergent Alternatives to Belief Propagation. *Neural Computation.* 2002; 14(7):1691–1722. <https://doi.org/10.1162/08997660260028674> PMID: [12079552](https://pubmed.ncbi.nlm.nih.gov/12079552/)

67. Kahneman D. Maps of Bounded Rationality: A Perspective on Intuitive Judgement. In: Frangsmyr T, editor. Nobel prizes, presentations, biographies, & lectures. Stockholm, Sweden: Almqvist & Wiksell; 2002. p. 416–499.
68. von Neumann J, Morgenstern O. Theory of Games and Economic Behavior. Princeton, NJ, USA: Princeton University Press; 1944.
69. Whittle P. Risk-sensitive optimal control. Chichester New York: Wiley; 1990.
70. Grau-Moya J, Leibfried F, Genewein T, Braun DA. Planning with Information-Processing Constraints and Model Uncertainty in Markov Decision Processes. In: Machine Learning and Knowledge Discovery in Databases. Springer International Publishing; 2016. p. 475–491.
71. Gottwald S, Braun DA. Systems of bounded rational agents with information-theoretic constraints. *Neural Computation*. 2019; 31(2):440–476. https://doi.org/10.1162/neco_a_01153 PMID: 30576612
72. Simon HA. A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*. 1955; 69(1):99–118. <https://doi.org/10.2307/1884852>
73. Marshall AW, Olkin I, Arnold BC. Inequalities: Theory of Majorization and Its Applications. 2nd ed. Springer New York; 2011.
74. Gottwald S, Braun DA. Bounded Rational Decision-Making from Elementary Computations That Reduce Uncertainty. *Entropy*. 2019; 21(4). <https://doi.org/10.3390/e21040375>
75. Ergin H, Sarver T. A Unique Costly Contemplation Representation. *Econometrica*. 2010; 78(4):1285–1339. <https://doi.org/10.3982/ECTA7801>
76. Todorov E. Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences*. 2009; 106(28):11478–11483. <https://doi.org/10.1073/pnas.0710743106> PMID: 19574462
77. Kappen HJ, Gómez V, Opper M. Optimal control as a graphical model inference problem. *Machine Learning*. 2012; 87(2):159–182. <https://doi.org/10.1007/s10994-012-5278-7>
78. Binz M, Gershman SJ, Schulz E, Endres D. Heuristics From Bounded Meta-Learned Inference. 2020;
79. Wolpert DH. The stochastic thermodynamics of computation. *Journal of Physics A: Mathematical and Theoretical*. 2019; 52(19):193001. <https://doi.org/10.1088/1751-8121/ab0850>
80. Miller GA. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*. 1956; 63(2):81–97. <https://doi.org/10.1037/h0043158> PMID: 13310704
81. Garner WR. Uncertainty and structure as psychological concepts. Wiley; 1962.
82. MacRae AW. Channel capacity in absolute judgment tasks: An artifact of information bias? *Psychological Bulletin*. 1970; 73(2):112–121. <https://doi.org/10.1037/h0028513>
83. Tatikonda S, Mitter S. Control Under Communication Constraints. *IEEE Transactions on Automatic Control*. 2004; 49(7):1056–1068. <https://doi.org/10.1109/TAC.2004.831187>
84. Harsha P, Jain R, McAllester D, Radhakrishnan J. The Communication Complexity of Correlation. *IEEE Transactions on Information Theory*. 2010; 56(1):438–449. <https://doi.org/10.1109/TIT.2009.2034824>
85. Genewein T, Leibfried F, Grau-Moya J, Braun DA. Bounded Rationality, Abstraction, and Hierarchical Decision-Making: An Information-Theoretic Optimality Principle. *Frontiers in Robotics and AI*. 2015; 2. <https://doi.org/10.3389/frobt.2015.00027>
86. Csiszár I. Axiomatic Characterizations of Information Measures. *Entropy*. 2008; 10(3):261–273. <https://doi.org/10.3390/e10030261>
87. Russell SJ, Subramanian D. Provably Bounded-optimal Agents. *Journal of Artificial Intelligence Research*. 1995; 2(1):575–609. <https://doi.org/10.1613/jair.133>
88. Gigerenzer G, Selten R. Bounded Rationality: The Adaptive Toolbox. MIT Press: Cambridge, MA, USA; 2001.
89. Ortega PA, Braun DA. Generalized Thompson sampling for sequential decision-making and causal inference. *Complex Adaptive Systems Modeling*. 2014; 2(1):2. <https://doi.org/10.1186/2194-3206-2-2>
90. Shannon CE. A Mathematical Theory of Communication. *The Bell System Technical Journal*. 1948; 27:379–656. <https://doi.org/10.1002/j.1538-7305.1948.tb00917.x>
91. Friston KJ, Kilner J, Harrison LM. A free energy principle for the brain. *Journal of Physiology-Paris*. 2006; 100:70–87. <https://doi.org/10.1016/j.jphysparis.2006.10.001> PMID: 17097864
92. Gershman SJ. What does the free energy principle tell us about the brain. *Neurons, Behavior, Data Analysis, and Theory*. 2019;
93. Wiener N. Cybernetics: Or Control and Communication in the Animal and the Machine. John Wiley; 1948.

94. Ashby W. Design for a Brain: The Origin of Adaptive Behavior. Springer Netherlands; 1960.
95. Powers WT. Behavior: The Control of Perception. Chicago, IL: Aldine; 1973.
96. Cisek P. Beyond the computer metaphor: behaviour as interaction. *Journal of Consciousness Studies*. 1999; 6(11-12):125–142.
97. Friston K. Life as we know it. *Journal of The Royal Society Interface*. 2013; 10(86):20130475. <https://doi.org/10.1098/rsif.2013.0475>
98. Corcoran AW, Hohwy J. Allostasis, interoception, and the free energy principle: Feeling our way forward. Oxford University Press; 2018.
99. Friston K, FitzGerald T, Rigoli F, Schwartenbeck P, O'Doherty J, Pezzulo G. Active inference and learning. *Neuroscience & Biobehavioral Reviews*. 2016; 68:862–879. <https://doi.org/10.1016/j.neubiorev.2016.06.022>
100. Friston KJ, FitzGerald THB, Rigoli F, Schwartenbeck P, Pezzulo G. Active Inference: A Process Theory. *Neural Computation*. 2017; 29:1–49. https://doi.org/10.1162/NECO_a_00912 PMID: 27870614
101. Schwöbel S, Kiebel S, Marković D. Active Inference, Belief Propagation, and the Bethe Approximation. *Neural Computation*. 2018; 30(9):2530–2567. https://doi.org/10.1162/neco_a_01108 PMID: 29949461
102. Parr T, Markovic D, Kiebel SJ, Friston KJ. Neuronal message passing using Mean-field, Bethe, and Marginal approximations. *Scientific Reports*. 2019; 9(1):1889. <https://doi.org/10.1038/s41598-018-38246-3> PMID: 30760782
103. Kikuchi R. A Theory of Cooperative Phenomena. *Physical Review*. 1951; 81(6):988–1003. <https://doi.org/10.1103/PhysRev.81.988>
104. Jeffrey RC. The Logic of Decision. 1st ed. University of Chicago Press; 1965.
105. Toussaint M, Storkey A. Probabilistic Inference for Solving Discrete and Continuous State Markov Decision Processes. In: Proceedings of the 23rd International Conference on Machine Learning. ICML'06. New York, NY, USA: Association for Computing Machinery; 2006. p. 945–952.
106. Todorov E. General duality between optimal control and estimation. In: 2008 47th IEEE Conference on Decision and Control. IEEE; 2008.
107. Levine S. Reinforcement Learning and Control as Probabilistic Inference: Tutorial and Review. arXiv:180500909. 2018;.
108. O'Donoghue B, Osband I, Ionescu C. Making Sense of Reinforcement Learning and Probabilistic Inference. In: International Conference on Learning Representations. ICLR'20; 2020.
109. Toussaint M. Robot trajectory optimization using approximate inference. In: Proceedings of the 26th Annual International Conference on Machine Learning—ICML'09. ACM Press; 2009. <https://doi.org/10.1145%2F1553374.1553508>
110. Ziebart BD. Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy. Carnegie Mellon University; 2010.
111. Tenenbaum JB, Griffiths TL. Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*. 2001; 24(4):629–640. <https://doi.org/10.1017/S0140525X01000061> PMID: 12048947
112. Gershman SJ, Daw ND. Perception, Action and Utility: The Tangled Skein. In: Principles of Brain Dynamics. MIT Press; 2012.
113. Dayan P, Hinton GE. Using Expectation-Maximization for Reinforcement Learning. *Neural Computation*. 1997; 9(2):271–278. <https://doi.org/10.1162/neco.1997.9.2.271>
114. Biehl M, Pollock FA, Kanai R. A technical critique of the free energy principle as presented in “Life as we know it” and related works. arXiv:200106408. 2020;.
115. Friston K, Costa LD, Parr T. Some interesting observations on the free energy principle. arXiv:200204501. 2020;.
116. Grünwald P. The Minimum Description Length Principle. Cambridge, Mass: MIT Press; 2007.
117. Schwartenbeck P, FitzGerald THB, Mathys C, Dolan R, Friston K. The Dopaminergic Midbrain Encodes the Expected Certainty about Desired Outcomes. *Cerebral cortex (New York, NY)*. 2015; 25(10):3434–3445. <https://doi.org/10.1093/cercor/bhu159> PMID: 25056572
118. Friston KJ, Parr T, de Vries B. The graphical brain: Belief propagation and active inference. *Network Neuroscience*. 2017; 1(4):381–414. https://doi.org/10.1162/NETN_a_00018 PMID: 29417960
119. Parr T, Benrimoh DA, Vincent P, Friston KJ. Precision and False Perceptual Inference. *Frontiers in Integrative Neuroscience*. 2018; 12:39. <https://doi.org/10.3389/fnint.2018.00039> PMID: 30294264
120. Rao RPN, Ballard DH. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*. 1999; 2(1):79–87. <https://doi.org/10.1038/4580> PMID: 10195184

121. Aitchison L, Lengyel M. With or without you: predictive coding and Bayesian inference in the brain. *Current Opinion in Neurobiology*. 2017; 46:219–227. <https://doi.org/10.1016/j.conb.2017.08.010> PMID: [28942084](https://pubmed.ncbi.nlm.nih.gov/28942084/)
122. Hohwy J. Self-supervision, normativity and the free energy principle. *Synthese*. 2020; <https://doi.org/10.1007/s11229-020-02622-2>
123. Alcock J. *Animal behavior: an evolutionary approach*. Sinauer Associates; 1993.
124. Tinbergen N. On aims and methods of Ethology. *Zeitschrift für Tierpsychologie*. 1963; 20:410–433. <https://doi.org/10.1111/j.1439-0310.1963.tb01161.x>