



Active Inference and Epistemic Value in Graphical Models

Thijs van de Laar^{1*}, Magnus Koudahl^{1,2}, Bart van Erp¹ and Bert de Vries^{1,3}

¹Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, Netherlands, ²Nested Minds Network Ltd., Liverpool, United Kingdom, ³GN Hearing Benelux BV, Eindhoven, Netherlands

OPEN ACCESS

Edited by:

Daniel Polani,
University of Hertfordshire,
United Kingdom

Reviewed by:

Akira Taniguchi,
Ritsumeikan University, Japan
Tim Verbelen,
Ghent University, Belgium

*Correspondence:

Thijs van de Laar
t.w.v.d.laar@tue.nl

Specialty section:

This article was submitted to
Computational Intelligence in Robotics,
a section of the journal
Frontiers in Robotics and AI

Received: 13 October 2021

Accepted: 27 January 2022

Published: 06 April 2022

Citation:

van de Laar T, Koudahl M, van Erp B
and de Vries B (2022) Active Inference
and Epistemic Value in
Graphical Models.
Front. Robot. AI 9:794464.
doi: 10.3389/frobt.2022.794464

The Free Energy Principle (FEP) postulates that biological agents perceive and interact with their environment in order to minimize a Variational Free Energy (VFE) with respect to a generative model of their environment. The inference of a policy (future control sequence) according to the FEP is known as Active Inference (AIF). The AIF literature describes multiple VFE objectives for policy planning that lead to epistemic (information-seeking) behavior. However, most objectives have limited modeling flexibility. This paper approaches epistemic behavior from a constrained Bethe Free Energy (CBFE) perspective. Crucially, variational optimization of the CBFE can be expressed in terms of message passing on free-form generative models. The key intuition behind the CBFE is that we impose a point-mass constraint on predicted outcomes, which explicitly encodes the assumption that the agent will make observations in the future. We interpret the CBFE objective in terms of its constituent behavioral drives. We then illustrate resulting behavior of the CBFE by planning and interacting with a simulated T-maze environment. Simulations for the T-maze task illustrate how the CBFE agent exhibits an epistemic drive, and actively plans ahead to account for the impact of predicted outcomes. Compared to an EFE agent, the CBFE agent incurs expected reward in significantly more environmental scenarios. We conclude that CBFE optimization by message passing suggests a general mechanism for epistemic-aware AIF in free-form generative models.

Keywords: free energy principle, active inference, variational optimization, constrained bethe free energy, message passing

Abbreviations: AIF, Active Inference; BFE, Bethe Free Energy; CBFE, Constrained Bethe Free Energy; EM, Expectation Maximization; FFG, Forney-style Factor Graph; FEP, Free Energy Principle; GM, Generative Model; MAP, Maximum A-Posteriori; VFE, Variational Free Energy; EFE, Expected Free Energy; VMP, Variational Message Passing; SP, Sum-Product.

1 INTRODUCTION

Free energy can be considered as a central concept in the natural sciences. Many natural laws can be derived through the principle of least action,¹ which rests on variational methods to minimize a path integral of free energy over time (Caticha, 2012). In neuroscience, an application of the least action principle to biological behavior is formalized as the Free Energy Principle (Friston et al., 2006). The Free Energy Principle (FEP) postulates that biological agents perceive and interact with their environment in order to minimize a Variational Free Energy (VFE) that is defined with respect to a model of their environment.

Under the FEP, perception relates to the process of hidden state estimation, where the agent tries to infer hidden causes of its sensory observations; and action (intervention) relates to a process where the agent actively tries to influence its (predicted) future observations by manipulating the external environment. Because the future is unobserved (by definition), the agent includes prior beliefs² about desired outcomes in its model and infers a policy that prescribes a sequence of future controls.³ The corollary of the FEP that includes action is referred to as Active Inference (AIF) (Friston et al., 2010).

The AIF literature describes multiple Free Energy (FE) objectives for policy planning, e.g., the Expected FE (Friston et al., 2015), Generalized FE (Parr and Friston, 2019) and Predicted (Bethe) FE (Schwöbel et al., 2018) (among others, see e.g. (Tschantz et al., 2020b; Hafner et al., 2020; Sajid et al., 2021)). Traditionally, the Expected Free Energy (EFE) is evaluated for a selection of policies, and a posterior distribution over policies is constructed from the corresponding EFEs. The EFE is designed to balance epistemic (knowledge seeking) and extrinsic (goal seeking) behavior. The active policy (the sequence of future controls to be executed in the environment) is then selected from this policy posterior (Friston et al., 2015).

Several authors have attempted to formulate minimization of the EFE by message passing on factor graphs (de Vries and Friston, 2017; Parr et al., 2019; Parr and Friston, 2019; Champion et al., 2021). These formulations evaluate the EFE objective with the use of a message passing scheme. In this paper we revisit this problem and compare the EFE approach with the message passing interpretation of the variational optimization of a Bethe Free Energy (BFE) (Pearl, 1988; Caticha, 2004; Yedidia et al., 2005). However, the BFE is known to lack epistemic (information-seeking) qualities, and resulting BFE AIF agents therefore do not pro-actively seek informative states (Schwöbel et al., 2018).

As a solution to the lack of epistemic qualities of the BFE, in this paper we approach epistemic behavior from a *Constrained*

BFE (CBFE) perspective (Şenöz et al., 2021). We illustrate how optimization of a point-mass constrained BFE objective instigates self-evidencing behavior. Crucially, variational optimization of the CBFE can be expressed in terms of message passing on a graphical representation of the underlying generative model (GM) (Dauwels, 2007; Cox et al., 2019), without modification of the GM itself. The contributions of this paper are as follows:

- We formulate the CBFE as an objective for epistemic-aware active inference (**Section 2.6**) that can be interpreted as message passing on a GM (**Section 6**);
- We interpret the constituent terms of the CBFE objective as drivers for behavior (**Section 4**);
- We illustrate our interpretation of the CBFE by planning and interacting with a simulated T-maze environment (**Section 5**).
- Simulations show that the CBFE agent plans epistemic policies multiple time-steps ahead (**Section 6.2**), and accrues reward for a significantly larger set of scenarios than the EFE (**Section 7**).

The main advantage of AIF with the CBFE objective, is that it allows inference to be fully automated by message passing, while retaining the epistemic qualities of the EFE. Automated message passing absolves the need for manual derivations and removes computational barriers in scaling AIF to more demanding settings (van de Laar, 2018).

2 PROBLEM STATEMENT

In this section we will introduce the free energy objectives as used throughout the paper. We start by introducing the Variational Free Energy (VFE), and explain how a VFE can be employed in an AIF context for perception and policy planning. We then introduce the Expected Free Energy (EFE) as a variational objective that is explicitly designed to yield epistemic behavior in AIF agents, but also note that the EFE definition limits itself to (hierarchical) state-space models. We then introduce the Bethe Free Energy (BFE), and argue that the BFE allows for convenient optimization on free-form models by message passing, but note that the BFE lacks information-seeking qualities. We conclude this section by introducing the Constrained Bethe Free Energy (CBFE), which equips the BFE with information-seeking qualities on free-form models through additional constraints on the variational distribution.

Table 1 summarizes notational conventions throughout the paper.

2.1 Variational Free Energy

The Variational Free Energy (VFE) is a principled metric in physics, where a time-integral over free energy is known as the action functional. Many natural laws can be derived from the principle of least action, where the action functional is minimized with the use of variational calculus (Caticha, 2012; Lanczos, 2012).

¹In this context, “action” refers to the path integral, and is distinct from “action” in the context of an intervention.

²We will use “belief” and “distribution” interchangeably.

³We use “controls” to refer to quantities in the generative model, and “actions” (in the intervention sense) to refer to quantities in the external environment.

TABLE 1 | Summary of notational conventions throughout the paper.

Notation	Def	Explanation
\mathbf{s}		Collection of (arbitrary) model variables
s_j		Individual model variable with index $j \in \mathcal{S}$
$f(\mathbf{s})$	(Eq. 1)	Factorized model of variables \mathbf{s}
$f_a(\mathbf{s}_a)$	(Eq. 1)	Factor (conditional or prior probability distribution) with argument variables \mathbf{s}_a and index $a \in \mathcal{F}$
$q(\mathbf{s})$	(Eq. 2)	Variational distribution of (latent) variables \mathbf{s}
$U_{q(\mathbf{s})}[f(\mathbf{s})]$	(Eq. 2)	Average energy
$H[q(\mathbf{s})]$	(Eq. 2)	Entropy
$F[q]$	(Eq. 2)	Variational Free Energy
y_k, x_k, u_k		Observation, state and control variable (at time k) respectively
\hat{y}_k		Specific realization for observation or unobserved future (predicted) outcome
\hat{u}_k		Specific control realization
$\rho(y_k, x_k x_{k-1}, u_k)$	(Eq. 7)	Generative Model engine
$\rho(y_k x_k)$	(Eq. 7)	Observation model
$\rho(x_k x_{k-1}, u_k)$	(Eq. 7)	Transition model
$\rho(x_{t-1})$	(Eq. 8)	State prior
$\mathbf{y}, \mathbf{x}, \mathbf{u}$		Sequence of future observation variables $\mathbf{y}_{t:t+T-1}$, state variables $\mathbf{x}_{t-1:t+T-1}$ and control variables $\mathbf{u}_{t:t+T-1}$, respectively
$\hat{\mathbf{u}}^j$		Policy (sequence of specific future controls), $\hat{\mathbf{u}}^j \in \mathcal{C}$, where index j is usually omitted
$F^*(\hat{\mathbf{u}}^j)$	(Eq. 12)	Optimized Variational Free Energy for policy $\hat{\mathbf{u}}^j$
$\hat{\mathbf{u}}^*$	(Eq. 13)	Optimal policy $\hat{\mathbf{u}}^* \in \mathcal{C}$
$G[q; \hat{\mathbf{u}}^j]$	(Eq. 14)	Expected Free Energy (EFE)
$\rho(\mathbf{y} \mathbf{x})$	(Eq. 16a)	Aggregate observation model
$\rho(\mathbf{x} \mathbf{u})$	(Eq. 16b)	Aggregate state transition model, including state prior
$\bar{\rho}(\mathbf{y})$	(Eq. 18)	Goal prior for sequence of future observation variables
$\bar{\rho}(y_k)$	(Eq. 18)	Goal prior for observation variable at time k
$f(\mathbf{y}, \mathbf{x} \mathbf{u})$	(Eq. 19), (Eq. 25)	Factorized model of future variables (at time t), for EFE and (C)BFE respectively
$B[q]$	(Eq. 23)	Bethe Free Energy (BFE)
$H_B[q]$	(Eq. 24)	Bethe entropy
$B[q; \hat{\mathbf{u}}^j]$	(Eq. 26)	Bethe Free Energy of future model under policy $\hat{\mathbf{u}}^j$
$B[q; \hat{\mathbf{u}}^j, \hat{\mathbf{y}}]$	(Eq. 28)	Constrained Bethe Free Energy (CBFE) of future model under policy $\hat{\mathbf{u}}^j$ and predicted outcomes $\hat{\mathbf{y}}$

The VFE is defined with respect to a factorized generative model (GM). We consider a GM $f(\mathbf{s})$ with factors $\{f_a | a \in \mathcal{F}\}$ and variables $\{s_i | i \in \mathcal{S}\}$ that factorizes according to

$$f(\mathbf{s}) = \prod_{a \in \mathcal{F}} f_a(\mathbf{s}_a), \tag{1}$$

where \mathbf{s}_a collects the argument variables of the factors f_a . As a notational convention, we write collections and sequences in bold script. In the model factorization of (Eq. 1), the factors f_a would correspond with the prior and conditional probability distributions that define the GM. The VFE is then defined as a functional of an (approximate) posterior $q(\mathbf{s})$ over latent variables, as

$$F[q] = \mathbb{E}_{q(\mathbf{s})} \left[\log \frac{q(\mathbf{s})}{f(\mathbf{s})} \right] = U_{q(\mathbf{s})}[f(\mathbf{s})] - H[q(\mathbf{s})], \tag{2}$$

which consists of an average energy $U_{q(\mathbf{s})}[f(\mathbf{s})] = -\mathbb{E}_{q(\mathbf{s})}[\log f(\mathbf{s})]$ and an entropy $H[q(\mathbf{s})] = -\mathbb{E}_{q(\mathbf{s})}[\log q(\mathbf{s})]$.

Because the VFE is (usually) optimized with respect to the posterior q with the use of variational calculus (Yedidia et al., 2005), the posterior q is also referred to as the *variational distribution*. In this paper, we will strictly reserve the q notation for variational distributions.

We can relate the exact posterior belief with the model definition through a normalizing constant Z , as

$$p(\mathbf{s}) = \frac{f(\mathbf{s})}{Z}, \tag{3}$$

where

$$Z = \sum_{\mathbf{s}} f(\mathbf{s}). \tag{4}$$

Throughout this paper, summation can be replaced by integration in the case of continuous variables.

In a Bayesian context, the normalizer Z is commonly referred to as the marginal likelihood or evidence for model f . However, exact summation (marginalization) of (Eq. 4) over all variable realizations is often prohibitively difficult in practice, so that the evidence and exact posterior become unobtainable.

Substituting (3) in the VFE (2) expresses the VFE as an upper bound on the surprise, that is the negative log-model evidence, as

$$F[q] = \underbrace{\text{KL}[q(\mathbf{s}) \| p(\mathbf{s})]}_{\text{posterior divergence}} \underbrace{-\log Z}_{\text{surprise}}. \tag{5}$$

The marginalization problem of (Eq. 4) is thus converted to an optimization problem over q . After optimization,

$$q^* = \arg \min_q F[q], \tag{6}$$

the VFE approximates the surprise, and the optimal variational distribution becomes an approximation to the true posterior, $p(\mathbf{s}) \approx q^*(\mathbf{s})$.

Crucially, we are free to choose constraints on q such that the optimization becomes practically feasible, at the cost of an increased posterior divergence. One such approximation is the Bethe assumption, as we will see in **Section 2.5**.

2.2 Inference for Perception

We now return to the model definition of (Eq. 1). In the context of AIF, a Generative Model (GM) comprises of a probability distribution over states x_k , observations y_k and controls u_k at each time index k . We will use a hat to indicate specific variable realizations, i.e. \hat{y}_k for a specific outcome and \hat{u}_k for a specific control at time k . As a notational convention, we use k as an arbitrary time index, often used in the context of iterations, and we use t to indicate the current simulation time index.

We define a state-space model (Koller and Friedman, 2009) for the generative model engine, which represents our belief about how observations follow from a given control and previous state, as

$$p(y_k, x_k | x_{k-1}, u_k) = \underbrace{p(y_k | x_k)}_{\text{observation model}} \underbrace{p(x_k | x_{k-1}, u_k)}_{\text{transition model}}. \quad (7)$$

We use a prior belief about past states $p(x_{t-1})$ together with the generative model engine (Eq. 7) to define a generative model for perception

$$f(y_t, x_t, x_{t-1} | u_t) = p(x_{t-1}) p(y_t, x_t | x_{t-1}, u_t) = p(x_{t-1}) p(y_t | x_t) p(x_t | x_{t-1}, u_t), \quad (8)$$

which, after substitution in (Eq. 2), results in the VFE objective for perception,

$$F[q] = \mathbb{U}_{q(x_t, x_{t-1})} [f(\hat{y}_t, x_t, x_{t-1} | \hat{u}_t)] - H[q(x_t, x_{t-1})]. \quad (9)$$

At each time t , the process of perception then relates to inferring the optimal variational distribution $q^*(x_t, x_{t-1})$ about latent states, given the current action \hat{u}_t and resulting outcome \hat{y}_t . The resulting variational distribution can then be used as a prior for the next time-step, such that $p(x_t) \triangleq q^*(x_t) = \sum_{x_{t-1}} q^*(x_t, x_{t-1})$.

2.3 Inference for Planning

At each time t , *planning* is concerned with selecting optimal future controls by minimizing a Free Energy (FE) objective that is defined with respect to future variables. We write $\mathbf{y} = \mathbf{y}_{t:t+T-1}$, $\mathbf{x} = \mathbf{x}_{t-1:t+T-1}$, and $\mathbf{u} = \mathbf{u}_{t:t+T-1}$ as the sequences of future observations, states and controls respectively, for a fixed-time horizon of T time-steps ahead.

We will refer to a specific future control sequence $\hat{\mathbf{u}}$ as a *policy*. The optimal policy $\hat{\mathbf{u}}^*$ is then referred to as the *active* policy, where (local) optimality is indicated by an asterisk. Inference for planning then aims to select the optimal policy (in terms of FE) from a collection of candidate policies $\hat{\mathbf{u}}^j \in \mathcal{C}$, where \mathcal{C} represents the finite set of all (user-provided) candidate policies, and j the selected policy index.

When we view the candidate policy $\hat{\mathbf{u}}^j$ as a model selection variable, the problem of policy selection becomes equivalent to the problem of Bayesian model selection, where we wish to find a probabilistic model with the highest posterior probability among some given candidate models. When there is no prior preference

about models, the optimal model is the one with the highest marginal likelihood (evidence).

Given a model $f(\mathbf{y}, \mathbf{x} | \mathbf{u})$ of future observations and states given a future control sequence, we can express the marginal likelihood (evidence) for a specific policy choice, as

$$Z_j = \sum_{\mathbf{y}} \sum_{\mathbf{x}} f(\mathbf{y}, \mathbf{x} | \hat{\mathbf{u}}^j). \quad (10)$$

Using (Eq. 3), we can then relate the exact posterior belief with the variational distribution and the policy evidence, as

$$p(\mathbf{y}, \mathbf{x} | \hat{\mathbf{u}}^j) = \frac{f(\mathbf{y}, \mathbf{x} | \hat{\mathbf{u}}^j)}{Z_j}. \quad (11)$$

Under optimization of q , the minimal VFE then approximates the surprise (Eq. 5), as

$$F^*(\hat{\mathbf{u}}^j) = \min_q F[q; \hat{\mathbf{u}}^j] \geq -\log Z_j. \quad (12)$$

The optimal policy then minimizes the optimized VFE, as

$$\hat{\mathbf{u}}^* = \arg \min_{\hat{\mathbf{u}}^j \in \mathcal{C}} F^*(\hat{\mathbf{u}}^j). \quad (13)$$

In the following, we omit the explicit indexing of the policy on j for notational convenience, and simply write $\hat{\mathbf{u}}$ to represent a specific policy choice.

Because the VFE (2) involves expectations over the full joint variational distribution, it may become prohibitively expensive to compute for larger models. Therefore, additional assumptions and constraints on the VFE are often required. As a result, multiple free energy objectives for policy planning have been proposed in the literature, e.g., the Expected Free Energy (EFE) (Friston et al., 2015; Friston et al., 2021), the Free Energy of the Expected Future (Millidge et al., 2020b), the Generalized Free Energy (Parr and Friston, 2019), the Predicted (Bethe) Free Energy (Schwöbel et al., 2018), and marginal approximations (Parr et al., 2019).

2.4 Expected Free Energy

The Expected Free Energy (EFE) is an FE objective for planning that is explicitly constructed to elicit information-seeking behavior (Friston et al., 2015). Because future observations are (by definition) unknown, the EFE is defined in terms of an expectation that includes observation variables, as

$$G[q; \hat{\mathbf{u}}] = \mathbb{E}_{q(\mathbf{y}, \mathbf{x})} \left[\log \frac{q(\mathbf{x})}{f(\mathbf{y}, \mathbf{x} | \hat{\mathbf{u}})} \right]. \quad (14)$$

Construction of the (Markovian) model for the EFE starts by stringing together a state prior with the generative model engine of (Eq. 7) for future times, as

$$p(\mathbf{y}, \mathbf{x} | \mathbf{u}) = p(x_{t-1}) \prod_{k=t}^{t+T-1} p(y_k, x_k | x_{k-1}, u_k) = p(x_{t-1}) \prod_{k=t}^{t+T-1} p(y_k | x_k) p(x_k | x_{k-1}, u_k), \quad (15)$$

where the state prior $p(x_{t-1})$ follows from the perceptual process (Section 2.2). For notational convenience, we often group the

observation and state transition models (including the state prior), according to

$$p(\mathbf{y}|\mathbf{x}) = \prod_{k=t}^{t+T-1} p(y_k|x_k) \quad (16a)$$

$$p(\mathbf{x}|\mathbf{u}) = p(x_{t-1}) \prod_{k=t}^{t+T-1} p(x_k|x_{k-1}, u_k). \quad (16b)$$

From the future generative model engine (Eq. 15), the EFE defines a state posterior

$$p(\mathbf{x}|\mathbf{y}, \mathbf{u}) = \frac{p(\mathbf{y}, \mathbf{x}|\mathbf{u})}{\sum_{\mathbf{x}} p(\mathbf{y}, \mathbf{x}|\mathbf{u})}. \quad (17)$$

Note that our notation differs from (Friston et al., 2015), where posterior distributions are denoted by q . We strictly reserve the q notation for variational distributions.

We introduce goal priors $\tilde{p}(y_k)$ over observation variables. Goal priors encode prior beliefs about desired observations (or states) (Friston et al., 2013), and are annotated with a tilde to avoid confusion with the marginal distribution over the same variable. We then introduce a shorthand notation that aggregates (independent) goal priors for the future generative model, as

$$\tilde{p}(\mathbf{y}) = \prod_{k=t}^{t+T-1} \tilde{p}(y_k). \quad (18)$$

Together with the aggregated goal prior, the factorized model for the EFE is then constructed as

$$f(\mathbf{y}, \mathbf{x}|\mathbf{u}) = p(\mathbf{x}|\mathbf{y}, \mathbf{u}) \tilde{p}(\mathbf{y}). \quad (19)$$

There are several things of note about the model of (Eq. 19):

- The model includes variables that pertain to future time-points, $t \leq k \leq t + T - 1$. As a result, the future observation variables \mathbf{y} are latent;
- The model includes a state prior that is a result of inference for perception;
- The (informative) goal priors \tilde{p} introduce a bias in the model towards desired outcomes;
- Candidate policies will be given, as indicated by a conditioning on controls.

Upon substitution of (Eq. 19), the EFE (Eq. 14) factorizes into an epistemic and an extrinsic value term (Friston et al., 2015), as

$$G[q; \hat{\mathbf{u}}] = \underbrace{-\mathbb{E}_{q(\mathbf{y}, \mathbf{x})} \left[\log \frac{p(\mathbf{x}|\mathbf{y}, \hat{\mathbf{u}})}{q(\mathbf{x})} \right]}_{\text{epistemic value}} - \underbrace{\mathbb{E}_{q(\mathbf{y})} [\log \tilde{p}(\mathbf{y})]}_{\text{extrinsic value}}, \quad (20)$$

where the epistemic value relates to a mutual information between states and observations. This decomposition is often used to motivate the epistemic qualities of the EFE.

An alternative decomposition, in terms of ambiguity and observation risk, can be obtained under the assumptions $q(\mathbf{y}|\mathbf{x}) \approx p(\mathbf{y}|\mathbf{x})$ (approximation of the observation model), and $q(\mathbf{x}|\mathbf{y}) \approx p(\mathbf{x}|\mathbf{y}, \mathbf{u})$ (approximation of the exact posterior). These assumptions allow us to rewrite the exact relationship

$q(\mathbf{y}, \mathbf{x}) = q(\mathbf{y}|\mathbf{x}) q(\mathbf{x}) = q(\mathbf{x}|\mathbf{y}) q(\mathbf{y})$ in terms of the approximations $q(\mathbf{y}, \mathbf{x}) \approx p(\mathbf{y}|\mathbf{x}) q(\mathbf{x}) \approx p(\mathbf{x}|\mathbf{y}, \mathbf{u}) q(\mathbf{y})$. As a result, we obtain

$$G[q; \hat{\mathbf{u}}] \approx \underbrace{\mathbb{E}_{q(\mathbf{y}, \mathbf{x})} \left[\log \frac{q(\mathbf{y})}{p(\mathbf{y}|\mathbf{x}) \tilde{p}(\mathbf{y})} \right]}_{\text{ambiguity}} + \underbrace{\text{KL}[q(\mathbf{y}) \| \tilde{p}(\mathbf{y})]}_{\text{observation risk}}, \quad (21)$$

where $q(\mathbf{x})$ and $q(\mathbf{y})$ on the r.h.s are implicitly conditioned on $\hat{\mathbf{u}}$. This decomposition is often used to motivate the explorative (ambiguity reducing) qualities of the EFE.

In the current paper we evaluate the EFE in accordance with (Friston et al., 2015; Da Costa et al., 2020a), for which the procedure is detailed in **Supplementary Appendix SA**.

Although the EFE leads to epistemic behavior, it does not fit the general functional form of the VFE (Eq. 2), where the expectation and numerator define the same variational distribution. As a result, EFE minimization by message passing requires custom definitions, and limits itself to (hierarchical) state-space models. Furthermore, note that the EFE involves the state posterior $p(\mathbf{x}|\mathbf{y}, \mathbf{u})$ as part of its definition, which is technically a quantity that needs to be inferred. The EFE thus conflates the definition of the planning objective with the inference procedure for planning itself.

2.5 Bethe Free Energy

The Bethe Free Energy (BFE) defines a variational distribution that factorizes according to the Bethe assumption

$$q(\mathbf{s}) \triangleq \prod_{a \in \mathcal{F}} q_a(\mathbf{s}_a) \prod_{i \in \mathcal{S}} q_i(\mathbf{s}_i)^{1-d_i}, \quad (22)$$

where the degree d_i counts how many q_a 's contain s_i as an argument. After substituting the Bethe assumption (Eq. 22) in the VFE (Eq. 2), we obtain the BFE,

$$B[q] = \sum_{a \in \mathcal{F}} \mathbb{U}_{q_a(\mathbf{s}_a)}[f_a(\mathbf{s}_a)] - \sum_{a \in \mathcal{F}} \mathbb{H}[q_a(\mathbf{s}_a)] - \sum_{i \in \mathcal{S}} (1 - d_i) \mathbb{H}[q_i(\mathbf{s}_i)], \quad (23)$$

as a special case of the VFE. The entropy contributions are often summarized in the Bethe entropy, as

$$\mathbb{H}_B[q] = \sum_{a \in \mathcal{F}} \mathbb{H}[q_a(\mathbf{s}_a)] + \sum_{i \in \mathcal{S}} (1 - d_i) \mathbb{H}[q_i(\mathbf{s}_i)]. \quad (24)$$

Because the BFE fully factorizes into local contributions in \mathcal{F} and \mathcal{S} , it can be optimized by message passing on the generative model (Pearl, 1988; Yedidia et al., 2005; Wainwright and Jordan, 2008; Şenöz et al., 2021). In the context of AIF, the BFE for a model over future states is also referred to as the Predicted Free Energy (Schwöbel et al., 2018).

For a fixed time-horizon T , the factorized model for future states is constructed from the generative model engine and goal prior, as

$$\begin{aligned} f(\mathbf{y}, \mathbf{x}|\mathbf{u}) &= p(\mathbf{y}, \mathbf{x}|\mathbf{u}) \tilde{p}(\mathbf{y}) \\ &= p(x_{t-1}) \prod_{k=t}^{t+T-1} p(y_k, x_k|x_{k-1}, u_k) \tilde{p}(y_k) \\ &= p(x_{t-1}) \prod_{k=t}^{t+T-1} p(y_k|x_k) p(x_k|x_{k-1}, u_k) \tilde{p}(y_k). \end{aligned} \quad (25)$$

Because the generative model engine and goal priors introduce a simultaneous constraint on future observations, the model of (Eq. 25) represents a scaled probability distribution. The BFE of the future model under policy $\hat{\mathbf{u}}$ then becomes

$$B[q; \hat{\mathbf{u}}] = U_{q(y,x)}[f(y, \mathbf{x}|\hat{\mathbf{u}})] - H_B[q(y, \mathbf{x})]. \quad (26)$$

A major advantage of the BFE over the EFE as an objective for AIF is that message passing implementations can be automatically derived on free form models, thus greatly enhancing model flexibility. A drawback of the BFE, however, is that it lacks the epistemic qualities of the EFE (Schwöbel et al., 2018), see also Section 4.

2.6 Constrained Bethe Free Energy

The *Constrained* Bethe Free Energy (CBFE) that we propose in this paper combines the epistemic qualities of the EFE with the computational ease and model flexibility of the BFE. The CBFE can be derived from the BFE by imposing additional constraints on the variational distribution

$$\begin{aligned} q(y, \mathbf{x}) &\hat{=} q(\mathbf{x}) \delta(y - \hat{y}) \\ &= q(\mathbf{x}) \prod_{k=t}^{t+T-1} \delta(y_k - \hat{y}_k), \end{aligned} \quad (27)$$

where $\delta(\cdot)$ defines the appropriate (Kronecker or Dirac) delta function for the domain of the observation variable y_k (discrete or continuous). The point-mass (delta) constraints of the CBFE are motivated by the following key insight: although the future is unknown, we know that we will observe something in the future. However, because future outcomes are by definition unobserved, the \hat{y}_k encode potential outcomes that need to be optimized for.

For the model of (Eq. 25), the CBFE then becomes⁴

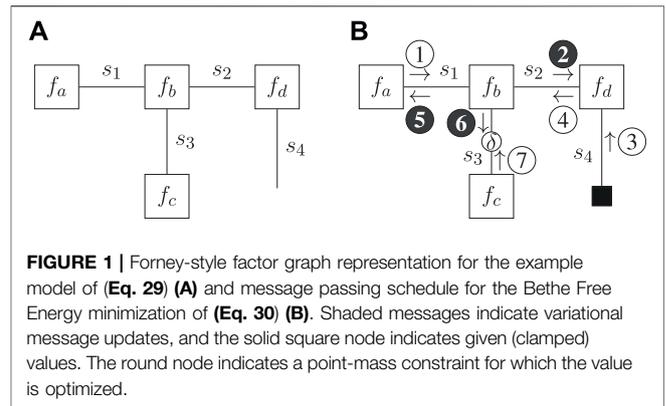
$$\begin{aligned} B[q; \hat{\mathbf{u}}, \hat{\mathbf{y}}] &= U_{q(\mathbf{x}) \delta(y-\hat{y})}[f(y, \mathbf{x}|\hat{\mathbf{u}})] - H_B[q(\mathbf{x}) \delta(y - \hat{y})] \\ &= U_{q(\mathbf{x})}[f(\hat{y}, \mathbf{x}|\hat{\mathbf{u}})] - H_B[q(\mathbf{x})]. \end{aligned} \quad (28)$$

The current paper investigates how point-mass constraints of the form (Eq. 27) affect epistemic behavior in AIF agents.

3 METHODS

To minimize the (Constrained) Bethe Free Energy, the current paper uses message passing on a Forney-style factor graph (FFG) representation (Forney, 2001) of the factorized model (Eq. 1). In an FFG, edges represent variables and nodes represent the functional relationships between variables (i.e. the prior and conditional probabilities).

Especially in a signal processing and control context, the FFG paradigm leads to convenient message passing formulations (Korl, 2005; Loeliger et al., 2007). Namely, inference can be described in terms of messages that summarize and propagate information across the FFG. The BFE is well-known for being the fundamental objective of the celebrated sum-product algorithm (Pearl, 1988), which has been formulated in terms of message passing on FFGs (Loeliger, 2004). Extensions of the sum-product algorithm to hybrid formulations, such as



variational message passing (VMP) (Dauwels, 2007) and expectation maximization (EM) (Dauwels et al., 2005) have also been formulated as message passing on FFGs. More recently, more general hybrid algorithms have been described in terms of message passing, see e.g. (Zhang et al., 2017; van de Laar et al., 2021). A comprehensive overview is provided in (Şenöz et al., 2021), where additional constraints, including point-mass constraints, are imposed on the BFE and optimized for by message passing on FFGs.

3.1 Forney-Style Factor Graph Example

Let us consider an example model (1) that factorizes according to

$$f(s_1, s_2, s_3, s_4) = f_a(s_1) f_b(s_1, s_2, s_3) f_c(s_3) f_d(s_2, s_4). \quad (29)$$

The FFG representation of (Eq. 29) is depicted in Figure 1A.

Now suppose we observe s_4 and introduce a point-mass constraint on s_3 . The variational distribution then factorizes as

$$\begin{aligned} q(s_1, s_2, s_3) &= q(s_1, s_2) q(s_3) \\ &= q(s_1, s_2) \delta(s_3 - \hat{s}_3), \end{aligned} \quad (30)$$

where s_4 is excluded from the variational distribution because it is observed and therefore no longer a latent variable. Substituting (Eq. 29) and (Eq. 30) in (Eq. 28) yields the CBFE as⁴

$$B[q; \hat{s}_3] = U_{q(s_1, s_2)}[f(s_1, s_2, \hat{s}_3, \hat{s}_4)] - H[q(s_1, s_2)], \quad (31)$$

where we directly substituted the observed value \hat{s}_4 into the factorized model.

In this paper we adhere to the notation in (Şenöz et al., 2021), and indicate point-mass constraints by an unshaded round node with an annotated δ on the corresponding edge of the FFG. A solid square node indicates a given value (e.g., an action, observed outcome or given parameter), whereas an unshaded round node indicates a point-mass constraint that is optimized for (e.g. a potential outcome). Unshaded messages indicate sum-product messages (Loeliger, 2004) and shaded messages indicate variational messages, as scheduled and computed in accordance with (Dauwels, 2007). The ForneyLab probabilistic

⁴For continuous variables we need to additionally assume that the entropy of a Dirac delta $H[\delta(\cdot)] = 0$ (Şenöz et al., 2021).

TABLE 2 | Optima for the individual terms of the CBEF decompositions (Eq. 32, Eq. 34). Each row varies one quantity (variable or function) in their respective term while other quantities remain fixed. Shaded cells indicate that the term (row) is not a function(al) of that specific optimization quantity (column).

Optimize	Vary			
	Fix	\hat{y}	\hat{u}	$q(x)$
max ex. val		\hat{y} that maximizes $\bar{\rho}(\hat{y})$		
min complexity	$q(x)$		$\hat{u} \in \mathcal{C}$ that renders $\rho(x \hat{u})$ closest to $q(x)$	$q(x) = \rho(x \hat{u})$
max confidence	$q(x)$	\hat{y} that maximizes expected outcomes ⁴ $\mathbb{E}_{q(x)}[\log \rho(\hat{y} x)]$		$q(x) = \delta(x - \hat{x})$ where \hat{x} maximizes the likelihood $\rho(\hat{y} x)$
Max intrinsic value	\hat{u}	\hat{y} that maximizes the evidence ⁴ $\rho(\hat{y} \hat{u})$	$\hat{u} \in \mathcal{C}$ that renders $\rho(y \hat{u})$ closest to $\delta(y - \hat{y})$	
min posterior divergence	$q(x), \hat{u}$	\hat{y} that renders $\rho(x \hat{y}, \hat{u})$ closest to $q(x)$	$\hat{u} \in \mathcal{C}$ that renders $\rho(x \hat{y}, \hat{u})$ closest to $q(x)$	$q(x) = \rho(x \hat{y}, \hat{u})$
	$q(x), \hat{y}$			
	\hat{y}, \hat{u}			

programming toolbox (Cox et al., 2019) implements an automated message passing scheduler and a lookup table of pre-derived message updates (Korl, 2005; van de Laar, 2019, **Supplementary Appendix SA**).

Variational optimization of (Eq. 31) then yields the (iterative) message passing schedule of **Figure 1B**, where \hat{s}_3 is initialized. After computation of the messages, the mode of the product between message ⑥ and ⑦ becomes the new value for \hat{s}_3 , and the schedule is repeated until convergence. The resulting optimization procedure then resembles an expectation maximization (EM) scheme where ⑥ acts as a likelihood and ⑦ as a prior (Dauwels et al., 2005), and where upon each iteration the value \hat{s}_3 is updated with the maximum a-posteriori (MAP) estimate.

4 VALUE DECOMPOSITIONS

In this section we further investigate the drivers for behavior of the (C)BEF. We assume that all variational distributions factorize according to the Bethe assumption (Eq. 22).

4.1 Confidence and Complexity

We substitute the model of (Eq. 25) in the CBEF definition of (Eq. 28) and combine to identify three terms, as

$$\begin{aligned}
 B[q; \hat{y}, \hat{u}] &= U_{q(x) \delta(y-\hat{y})}[p(y, x|\hat{u})] + U_{\delta(y-\hat{y})}[\tilde{p}(y)] \\
 &- H_B[q(x)\delta(y-\hat{y})] = U_{q(x) \delta(y-\hat{y})}[p(y|x)] - H_B[\delta(y-\hat{y})] \\
 &+ U_{q(x)}[p(x|\hat{u})] - H_B[q(x)] + U_{\delta(y-\hat{y})}[\tilde{p}(y)] \\
 &= \underbrace{\mathbb{E}_{q(x)}[\text{KL}[\delta(y-\hat{y})\|p(y|x)]]}_{\text{negative confidence}} + \underbrace{\text{KL}[q(x)\|p(x|\hat{u})]}_{\text{complexity}} - \underbrace{\log \tilde{p}(\hat{y})}_{\text{extrinsic value}}.
 \end{aligned}
 \tag{32}$$

Table 2 is provided as an overview, and summarizes the properties of the individual terms of (Eq. 32) under optimization.

The extrinsic value induces a preference for extrinsically rewarding future outcomes.

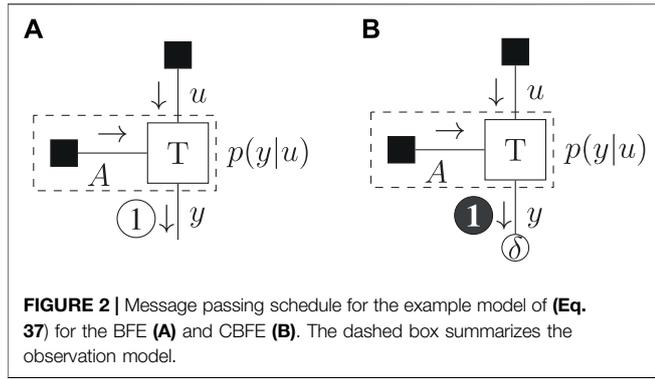
Minimizing complexity prefers policies that induce transitions that are in line with state beliefs, and (vice versa) prefers state beliefs that remain close the policy-induced state transitions.

The confidence expresses the expected difference (divergence) in information between the outcomes as predicted by the observation model and the most likely (expected) outcome. In other words, this term quantifies the information difference between predictions and absolute certainty about outcomes. While the negative confidence term could be interpreted as an ambiguity (deviation from certainty), we choose this alternative terminology to prevent confusion with the ambiguity as defined in (21).

Specifically, the ambiguity (Eq. 21) and negative confidence (Eq. 32) are both of the form⁴ $U_{q(y,x)}[p(y|x)] = -\mathbb{E}_{q(y,x)}[\log p(y|x)]$. However, where the ambiguity approximates $q(y, x) \approx p(y|x) q(x)$ (Eq. 21), the confidence defines $q(y, x) \triangleq q(x) \delta(y - \hat{y})$ (Eq. 27). As a result, the ambiguity explicitly accounts for the full spread of $p(y|x)$, whereas the confidence evaluates $p(y|x)$ at the expected maximum \hat{y} .

Maximizing confidence prefers outcomes that are in line with predictions, and simultaneously tries to maximize the precision of state beliefs (**Table 2**), see also (Friston and Penny, 2011, p. 2093). Note that all terms act in unison—the precision of state beliefs is simultaneously influenced by the complexity, which prevents the collapse of the state belief to a point-mass.

⁴For continuous variables we need to additionally assume that the entropy of a Dirac delta $H[\delta(\cdot)] = 0$ (Şenöz et al., 2021).



4.2 Intrinsic and Extrinsic Value

A second decomposition of the CBFE objective follows when we rewrite the factorized model of (Eq. 25) using the product rule, as

$$f(y, \mathbf{x}|\mathbf{u}) = p(y, \mathbf{x}|\mathbf{u}) \tilde{p}(y) = p(\mathbf{x}|\mathbf{y}, \mathbf{u}) p(\mathbf{y}|\mathbf{u}) \tilde{p}(y). \quad (33)$$

Substituting (Eq. 33) in the CBFE definition (Eq. 28) and combining terms, then yields

$$\begin{aligned} B[q; \hat{\mathbf{y}}, \hat{\mathbf{u}}] &= U_{q(\mathbf{x})} \delta_{(\mathbf{y}-\hat{\mathbf{y}})} [p(\mathbf{y}, \mathbf{x}|\hat{\mathbf{u}})] + U_{\delta_{(\mathbf{y}-\hat{\mathbf{y}})}} [\tilde{p}(y)] \\ &- H_B[q(\mathbf{x})\delta(\mathbf{y}-\hat{\mathbf{y}})] = U_{q(\mathbf{x})} [p(\mathbf{x}|\hat{\mathbf{y}}, \hat{\mathbf{u}})] - H_B[q(\mathbf{x})] \\ &+ U_{\delta_{(\mathbf{y}-\hat{\mathbf{y}})}} [p(\mathbf{y}|\hat{\mathbf{u}})] - H_B[\delta(\mathbf{y}-\hat{\mathbf{y}})] + U_{\delta_{(\mathbf{y}-\hat{\mathbf{y}})}} [\tilde{p}(y)] \\ &= \underbrace{\text{KL}[q(\mathbf{x})\|p(\mathbf{x}|\hat{\mathbf{y}}, \hat{\mathbf{u}})]}_{\text{posterior divergence}} + \underbrace{\text{KL}[\delta(\mathbf{y}-\hat{\mathbf{y}})\|p(\mathbf{y}|\hat{\mathbf{u}})]}_{\text{negative intrinsic value}} - \underbrace{\log \tilde{p}(\hat{\mathbf{y}})}_{\text{extrinsic value}}. \end{aligned} \quad (34)$$

Table 2 again summarizes the properties of the individual terms of (Eq. 34) under optimization.

The second term of (Eq. 34) expresses the difference (divergence) in information between the predicted outcomes and the point-mass constrained (expected) outcome. In contrast to the extrinsic value (third term), this term quantifies a (negative) intrinsic value that purely depends upon the agent’s intrinsic beliefs about the environment (state prior and generative model engine (Eq. 15)). Under optimization (Table 2), this term prefers policies that lead to precise predictions for the outcomes.

The posterior divergence (first term) is always non-negative and will diminish under optimization, which allows us to combine (Eq. 32) and (Eq. 34) into Eq. 4

$$\underbrace{\log p(\hat{\mathbf{y}}|\hat{\mathbf{u}})}_{\text{predicted log-evidence}} \geq \underbrace{\mathbb{E}_{q(\mathbf{x})}[\log p(\hat{\mathbf{y}}|\mathbf{x})]}_{\text{confidence}} - \underbrace{\text{KL}[q(\mathbf{x})\|p(\mathbf{x}|\hat{\mathbf{u}})]}_{\text{complexity}}. \quad (35)$$

Interestingly, (Eq. 35) tells us that the intrinsic value of (Eq. 34) relates to the model evidence as predicted under the policy and resulting expected outcomes. Inference for planning with the CBFE then attempts to make precise predictions for outcomes by maximizing predicted model evidence. In this view, the CBFE for planning can be considered (quite literally) as self-evidencing (Friston et al., 2010; Hohwy, 2016).

As a result of selected actions, environmental outcomes may still be surprising under current generative model assumptions. Inference for perception then subsequently corrects the generative model priors (Section 2.2), and the action-perception loop repeats (see also Algorithm 1). Epistemic qualities then emerge from this continual pursuit of evidence (Hohwy, 2021).

In short, we note a distinction in the interpretation of epistemic value between the EFE and the CBFE. In the EFE (Eq. 20), epistemic value is directly related with a mutual information term between states and outcomes. In the CBFE, the epistemic drive appears to result from a self-evidencing mechanism.

4.3 Bethe Free Energy Value Decomposition

The BFE does not permit an interpretation in terms of intrinsic value. When we substitute (Eq. 33) in the BFE definition of (Eq. 26) and combine terms, we obtain

$$\begin{aligned} B[q; \hat{\mathbf{u}}] &= U_{q(\mathbf{y}, \mathbf{x})} [p(\mathbf{y}, \mathbf{x}|\hat{\mathbf{u}})] + U_{q(\mathbf{y})} [\tilde{p}(y)] - H_B[q(\mathbf{y}, \mathbf{x})] \\ &= \text{KL}[q(\mathbf{y}, \mathbf{x})\|p(\mathbf{y}, \mathbf{x}|\hat{\mathbf{u}})] - \mathbb{E}_{q(\mathbf{y})}[\log \tilde{p}(y)] \\ &= \underbrace{\mathbb{E}_{q(\mathbf{y})}[\text{KL}[q(\mathbf{x}|\mathbf{y})\|p(\mathbf{x}|\mathbf{y}, \hat{\mathbf{u}})]]}_{\text{expected posterior divergence}} + \underbrace{\text{KL}[q(\mathbf{y})\|p(\mathbf{y}|\hat{\mathbf{u}})]}_{\text{predictive divergence}} \\ &- \underbrace{\mathbb{E}_{q(\mathbf{y})}[\log \tilde{p}(y)]}_{\text{expected extrinsic value}}. \end{aligned} \quad (36)$$

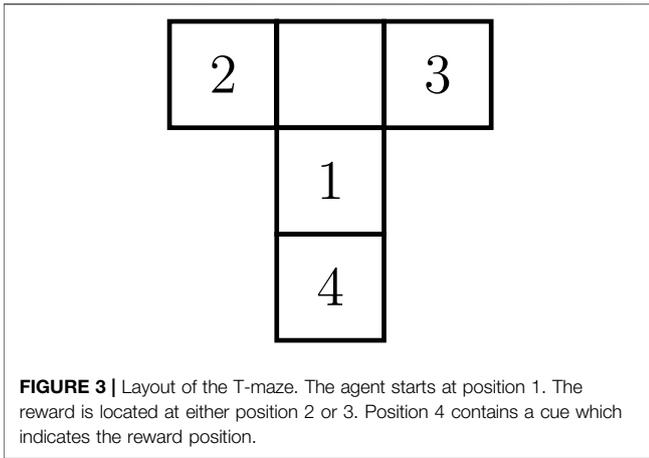
The intrinsic value term of (Eq. 34) has been replaced by a predictive divergence in (Eq. 36). This term expresses the difference (divergence) in information between the observations as predicted by the model under policy $\hat{\mathbf{u}}$, and variational distribution about outcomes $q(\mathbf{y})$. Under optimization of $q(\mathbf{x}|\mathbf{y})$, the posterior divergence will vanish for all \mathbf{y} . Without the additional point-mass constraint of (Eq. 27), the predictive divergence then no longer quantifies the information difference between uncertainty (predictive distribution) and certainty (predicted outcomes). As a result, the BFE lacks the self-evidencing qualities and resulting epistemic drive of the CBFE, as we will further illustrate in Section 4.4, 6.

4.4 Example Application

We illustrate our interpretation of (Eqs 34, 36) by a minimal example model. We consider a two-armed bandit, where an agent chooses between two levers, $u \in \{0, 1\}$. Each lever offers a distinct probability for observing an outcome $y \in \{0, 1\}$. Specifically, choosing $\hat{u} = 0$ will offer a 0.5 probability for observing $\hat{y} = 0$ (ignorant policy), whereas choosing $\hat{u} = 1$ will always lead to the observation $\hat{y} = 0$ (informative policy). We do not equip the agent with any external

TABLE 3 | Free energies (in bits) per policy for the example application.

Policy	BFE	CBFE
Ignorant ($\hat{u} = 0$)	0	1
Informative ($\hat{u} = 1$)	0	0



preference (there is no extrinsic reward). The agent’s factorized model then becomes

$$f(y|u) = p(y = a_i|u = a_j) = A_{ij}, \quad (37)$$

with $a = (0,1)^T$ and the conditional probability matrix

$$A = \begin{pmatrix} 0.5 & 1 \\ 0.5 & 0 \end{pmatrix}. \quad (38)$$

The BFE then follows as

$$B[q; \hat{u}] = U_{q(y)}[p(y|\hat{u})] - H[q(y)]. \quad (39)$$

The CBFE additionally constrains $q(y) = \delta(y - \hat{y})$, and as a result corresponds directly with the negative intrinsic value term of (Eq. 34), as⁴

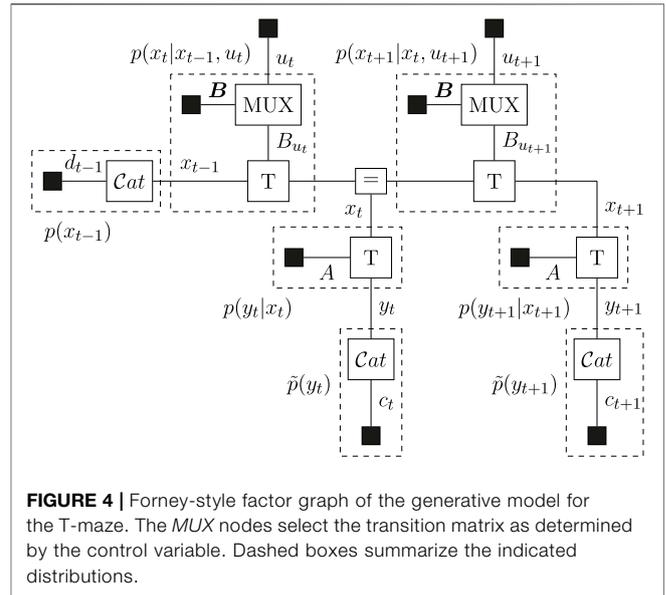
$$B(\hat{y}, \hat{u}) = U_{\delta(y-\hat{y})}[p(y|\hat{u})] - H[\delta(y - \hat{y})] = -\log p(\hat{y}|\hat{u}). \quad (40)$$

The FFG for the model definition of (Eq. 37) together with the resulting schedule for optimization of the (C)BFE is drawn in Figure 2.

The results of Table 3 show that the BFE does not distinguish between policies. The CBFE however penalizes the ignorant policy ($\hat{u} = 0$), which does not predict precise outcomes. This mechanism thus induces a preference for the informative policy ($\hat{u} = 1$), which does predict precise outcomes. In the following, we will further investigate this behavior in a less trivial setting.

5 EXPERIMENTAL SETTING

A classic experimental setting that investigates epistemic behavior is the T-maze task (Friston et al., 2015). The T-maze environment consists of four positions $\mathcal{P} = \{1, 2, 3, 4\}$, as drawn in Figure 3. The agent starts in position 1, and aims to obtain a reward that resides in either arm 2 or 3, $\mathcal{R} = \{2, 3\}$. The position of the reward



is unknown to the agent a priori, and once the agent enters one of the arms it remains there.

In order to learn the position of the reward, the agent first needs to move to position 4, where a cue indicates the reward position. At each position, the agent may observe one of four reward-related outcomes $\mathcal{O} = \{1, 2, 3, 4\}$:

1. The reward is indicated to reside at location two (left arm);
2. The reward is indicated to reside at location three (right arm);
3. The reward is obtained;
4. The reward is not obtained.

The key insight is that an epistemic policy would first inspect the cue at position 4 and then move to the indicated arm, whereas a purely goal directed agent would immediately move towards either of the potential goal positions instead of visiting the cue.

5.1 Generative Model Specification

We follow (Friston et al., 2015), and assume a generative model with discrete states x_k , observations y_k and controls u_k . The state $x_k \in \mathcal{P} \times \mathcal{R}$, represents the agent position at time k (four positions, Figure 3) combined with the reward position (two possibilities). The state vector thus comprises of eight possible realizations. The transition between states is affected by the control $u_k \in \mathcal{P}$, which encodes the agent’s attempted next position in the maze. The observation variables $y_k \in \mathcal{O} \times \mathcal{P}$ represent the agent position at time k (four positions) in combination with the reward-related outcome at that position (four possibilities).

The respective state prior, observation model, transition model and goal priors are defined as

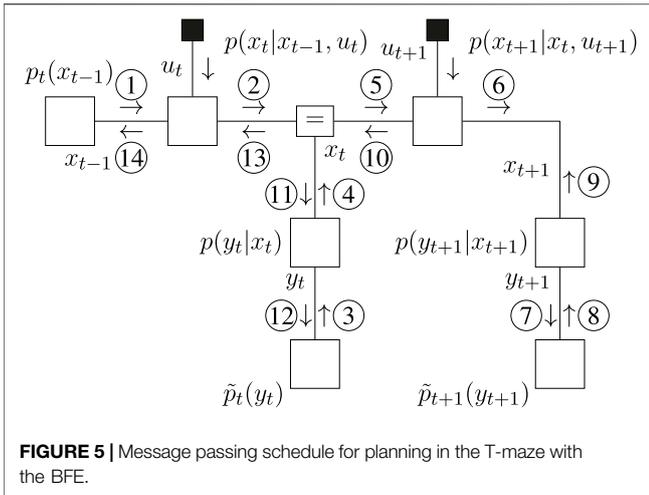


FIGURE 5 | Message passing schedule for planning in the T-maze with the BFE.

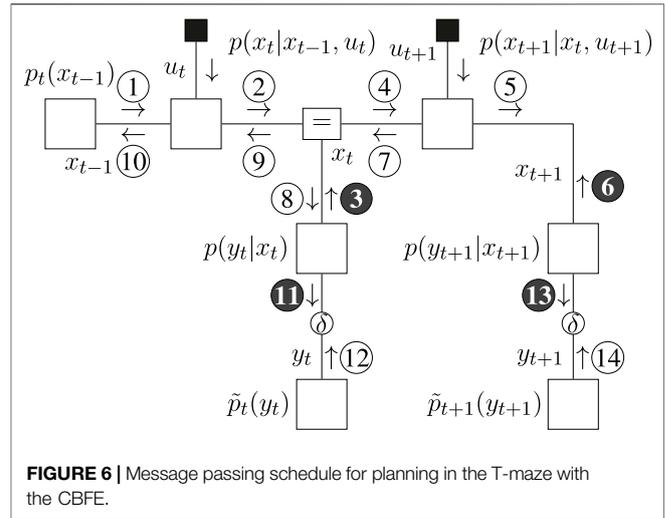


FIGURE 6 | Message passing schedule for planning in the T-maze with the CBF.

$$\begin{aligned}
 p(x_{t-1}) &= \text{Cat}(x_{t-1}|d_{t-1}) \\
 p(y_k = a_i|x_k = b_j) &= A_{ij} \\
 p(x_k = b_i|x_{k-1} = b_j, u_k = \hat{u}_k) &= (B_{\hat{u}_k})_{ij} \\
 \tilde{p}(y_k) &= \text{Cat}(y_k|c_k),
 \end{aligned}$$

where $a_i \in \mathcal{O} \times \mathcal{P}$, $b_i \in \mathcal{P} \times \mathcal{R}$, and $b_j \in \mathcal{P} \times \mathcal{R}$.

The agent plans two steps ahead ($T = 2$), for which the FFG is drawn in **Figure 4**.

5.2 Parameter Assignments

We start the simulation at $t = 1$, and assume that the initial position is known, namely we start at position 1 (**Figure 3**). However, the reward position is unknown a priori. This prior information is encoded by the initial state probability vector

$$d_0 = (1, 0, 0, 0)^T \otimes (0.5, 0.5)^T,$$

where \otimes denotes the Kronecker product.

The transition matrix $B_{\hat{u}_k}$ encodes the state transitions (from column-index to row-index), as

$$\begin{aligned}
 B_1 &= \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes I_2, & B_2 &= \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes I_2, \\
 B_3 &= \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes I_2, & B_4 &= \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix} \otimes I_2.
 \end{aligned}$$

The control affects the agent position, but not the reward position. Therefore, Kronecker products with the two-dimensional unit matrix I_2 ensure that the transitions are duplicated for both possible reward positions. Note that positions 2 and 3 (the reward arms) are attracting states, since none of the transition matrices allow a transition away from these positions. This means that although it is possible to propose any control at any time, not all controls will move the agent to its attempted position. We denote the collection of transition matrices by $\mathbf{B} = \{B_1, B_2, B_3, B_4\}$.

The observed outcome depends on the position of the agent. The position-dependent observation matrices specify how observations follow, given the current position of the agent (subscripts) and the reward position (columns), as

$$\begin{aligned}
 A_1 &= \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, & A_2 &= \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ \alpha & 1 - \alpha \\ 1 - \alpha & \alpha \end{pmatrix}, \\
 A_3 &= \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 - \alpha & \alpha \\ \alpha & 1 - \alpha \end{pmatrix}, & A_4 &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix},
 \end{aligned} \tag{41}$$

with reward probability α . The columns of these position-dependent observation matrices represent the two possibilities for the reward position. The position-dependent observation matrices combine into the complete block-diagonal, 16-by-8 observation matrix

$$A = A_1 \oplus A_2 \oplus A_3 \oplus A_4,$$

where \oplus denotes the direct sum (i.e. block-diagonal concatenation).

The goal prior depends upon the future time,

$$c_1 = (0.25, 0.25, 0.25, 0.25)^T \otimes (0.25, 0.25, 0.25, 0.25)^T \tag{42a}$$

$$c_k = \alpha \left((0, 0, c, -c)^T \otimes (1, 1, 1, 1)^T \right) \text{ for } k > 1, \tag{42b}$$

with reward utility c , and σ the soft-max function where $\sigma(s)_i = \frac{\exp(s_i)}{\sum_j \exp(s_j)}$. The flat prior c_1 encodes a lack of external preference at $t = 1$, while c_k for $k > 1$ encodes a preference for observing rewards at subsequent times. This effectively removes the goal prior for the first move ($t = 1$), while in subsequent moves the agent is rewarded for extrinsically rewarding states (van de Laar, 2018).

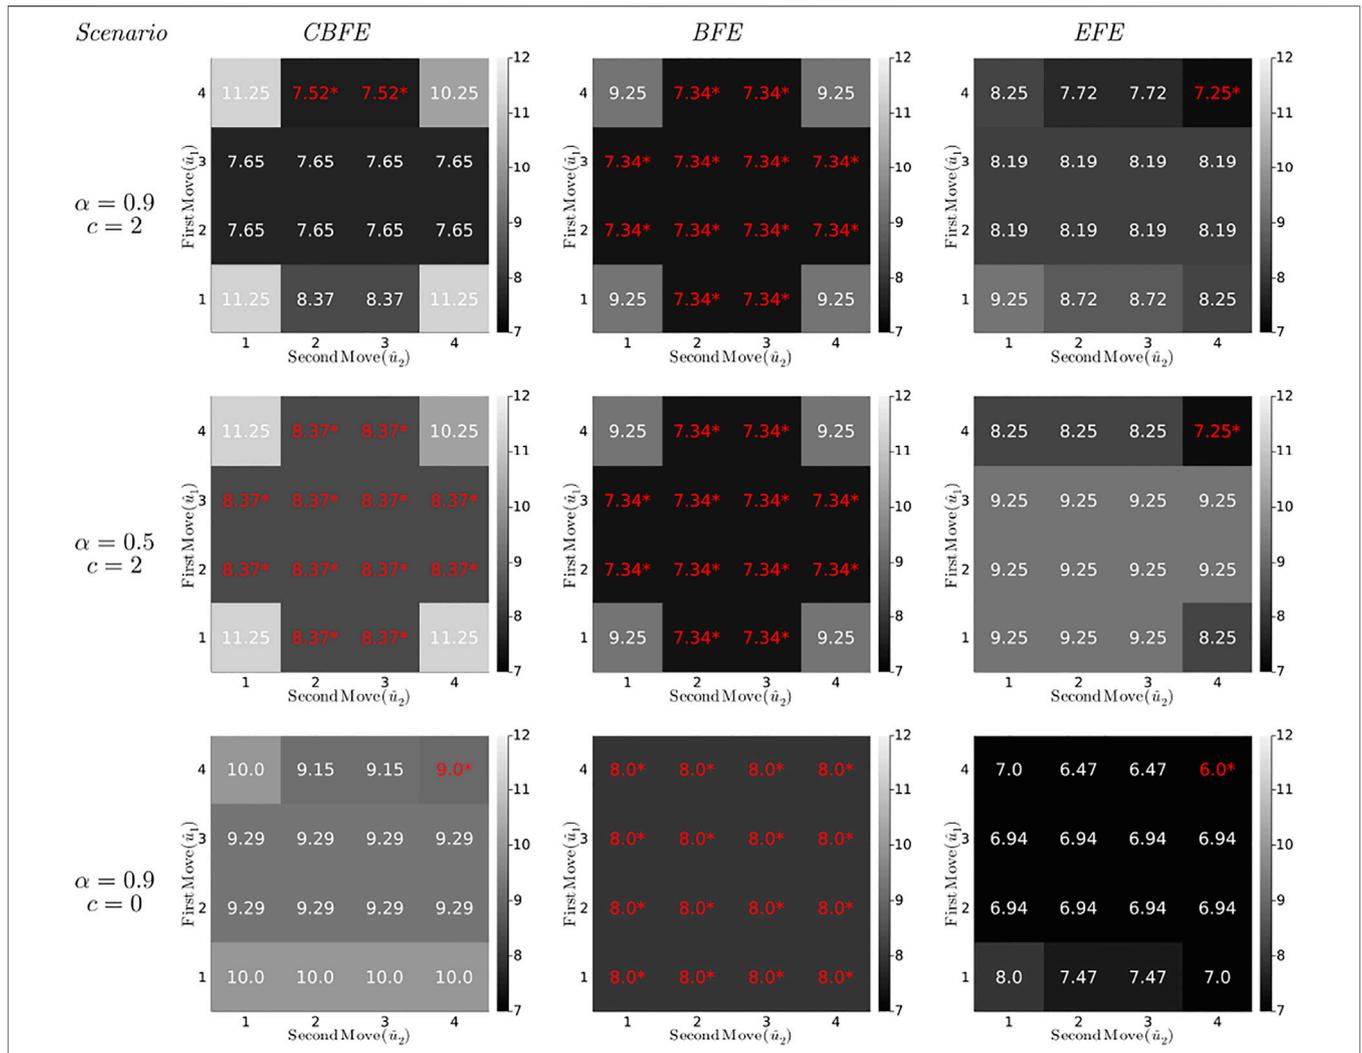


FIGURE 7 | (Constrained) Bethe Free Energies ((C)BFE) and Expected Free Energies (EFE) (in bits) for the T-maze policies under varying parameter settings. Each diagram plots the minimized free energy values for all possible policies (lookahead $T = 2$), with the first move on the vertical axis and the second move on the horizontal axis. For example, the cell in row 4, column 3 represents the policy $\hat{\mathbf{u}} = (4, 3)$, which first moves to position 4 and then to position 3. The values for the optimal policies $\hat{\mathbf{u}}^*$ are annotated red with an asterisk.

6 INFERENCE FOR PLANNING

In this simulation we compare the behavior of a CBFE agent to the behavior of a reference BFE agent (without point-mass constraints). We consider given policies $\hat{\mathbf{u}}$, and the optimal CBFE as a function of those policies

$$B_y^*(\hat{\mathbf{u}}) = \min_{q, \hat{\mathbf{y}}} B[q; \hat{\mathbf{y}}, \hat{\mathbf{u}}], \quad (43)$$

where the $\hat{\mathbf{y}}$ subscript indicates the explicit inclusion of point-mass constraints.

The unconstrained BFE represents the objective where the future observation variables $\hat{\mathbf{y}}$ are not point-mass constrained by their potential outcomes $\hat{\mathbf{y}}$. The unconstrained agent will

therefore optimize the joint belief over state and future observation variables rather than potential outcomes, as

$$B^*(\hat{\mathbf{u}}) = \min_q B[q; \hat{\mathbf{u}}]. \quad (44)$$

We will evaluate the BFE, CBFE and EFE for all sixteen ($T = 2$) possible candidate policies $\hat{\mathbf{u}} \in \mathcal{U} \times \mathcal{U}$. We consider several scenarios with varying reward probabilities α (Eq. 41) and reward utilities c (42).

In the current section we do not (yet) consider the interaction of the agent with the environment. In other words, actions from optimal policy $\hat{\mathbf{u}}^*$ are not (yet) executed; we are purely interested in the inference for planning itself, and the resulting free energy values as a function of the candidate policies (Eq. 43).

6.1 Message Passing Schedule for Planning

The message passing schedules for planning are drawn in **Figure 5** (BFE) and **Figure 6** (CBFE), where light messages are computed by sum-product (SP) message passing updates (Loeliger, 2004), and dark messages by variational message passing updates (Dauwels, 2007). An overview of message passing updates for discrete nodes can be found in (van de Laar, 2019, **Supplementary Appendix SA**).

For the CBFE, the posterior beliefs associated with the observation variables are constrained by point-mass (Dirac-delta) distributions, see (27), and the corresponding potential outcomes are optimized for. The message passing optimization scheme is derived from first principles in (Şenöz et al., 2021). In order to obtain a new value, e.g. \hat{y}_t , messages \square and \otimes are multiplied. The mode of the product then becomes the new value \hat{y}_t , which is used to construct the belief $q(y_t) = \delta(y_t - \hat{y}_t)$. The updated belief is subsequently used in the next iteration to compute \ominus . The resulting iterative expectation maximization (EM) procedure initializes values for all \hat{y}_k , and is performed using message passing according to (Dauwels et al., 2005). Interestingly, where optimization of the EFE is performed by a forward-only procedure (see **Supplementary Appendix SA**), optimization of the (C)BFE, as illustrated in **Figures 5, 6**, also includes a complete backward (smoothing) pass over the model.

6.2 Inference Results for Planning

Optimization of the (C)BFE by message passing is performed with ForneyLab^{5,6} version 0.11.4 (Cox et al., 2019). Free energies for planning, for three different agents and T-maze scenarios, are plotted in **Figure 7**. The distinct agents optimize the CBFE, BFE and EFE, respectively. We summarize the most important observations below.

The first column of diagrams in **Figure 7** shows the results for the CBFE agent, for varying scenarios.

- The first scenario for the CBFE agent (upper left diagram) imposes a likely reward ($\alpha = 0.9$) and positive reward utility ($c = 2$). In this scenario, the CBFE agent prefers the informative policies (4,2) and (4,3), where the agent seeks the cue in the first move and the reward in the second move. An epistemic (information seeking) agent would prefer these policies in this scenario.
- In the upper left diagram, note the lack of preference between position 2 and 3 in the second move. Because the policy is not yet executed (moves are only planned), the true reward location remains unknown. Therefore, both of these informative policies are on equal footing.

The second column of diagrams shows the results for the BFE agent.

- In every scenario, the BFE agent fails to distinguish between the majority of ignorant (first move to 1), informative (first move to 4) and greedy policies (first move to 2 or 3). These policy preferences do not correspond with the anticipated preferences of an epistemic agent.
- Comparing the BFE with the CBFE results, we observe that the point-mass constraint on potential outcomes induces a differentiation between ignorant, informative and greedy policies.
- More specifically, the third scenario (third row of diagrams) removes the extrinsic value of reward ($c = 0$). While the CBFE still differentiates between ignorant, informative and greedy policies, the BFE agent exhibits a total lack of preference.
- The second scenario (second row of diagrams) removes the value of information about the reward position ($\alpha = 0.5$). This scenario thus renders the cue worthless. The BFE agent appears unsusceptible to a change in the epistemic α parameter.

Taken together, these observations support the interpretation of the BFE as a purely extrinsically driven objective (**Section 4**).

The third column of diagrams produces the results for an EFE agent, as implemented in accordance with (Friston et al., 2015), see also **Supplementary Appendix SA**.

- In all scenarios, the EFE agent exhibits a consistent preference for the (4,4) policy. Compared to the CBFE agent, the EFE agent fails to plan ahead to obtain future reward after observing the cue.
- As we will see in **Section 7**, the EFE agent only infers a preference for a reward arm (position 2 or 3) after *execution* of the first move to the cue position. In contrast, the CBFE agent predicts the impact of information and plans accordingly.
- The second scenario (middle row) provides an informative cue, but removes the possibility to exploit that information ($\alpha = 0.5$). Interestingly, the EFE agent still moves to the cue position (for the sake of getting information), whereas the CBFE agent expresses ambivalence under an inoperable cue.

6.3 Results for CBFE Value Decomposition

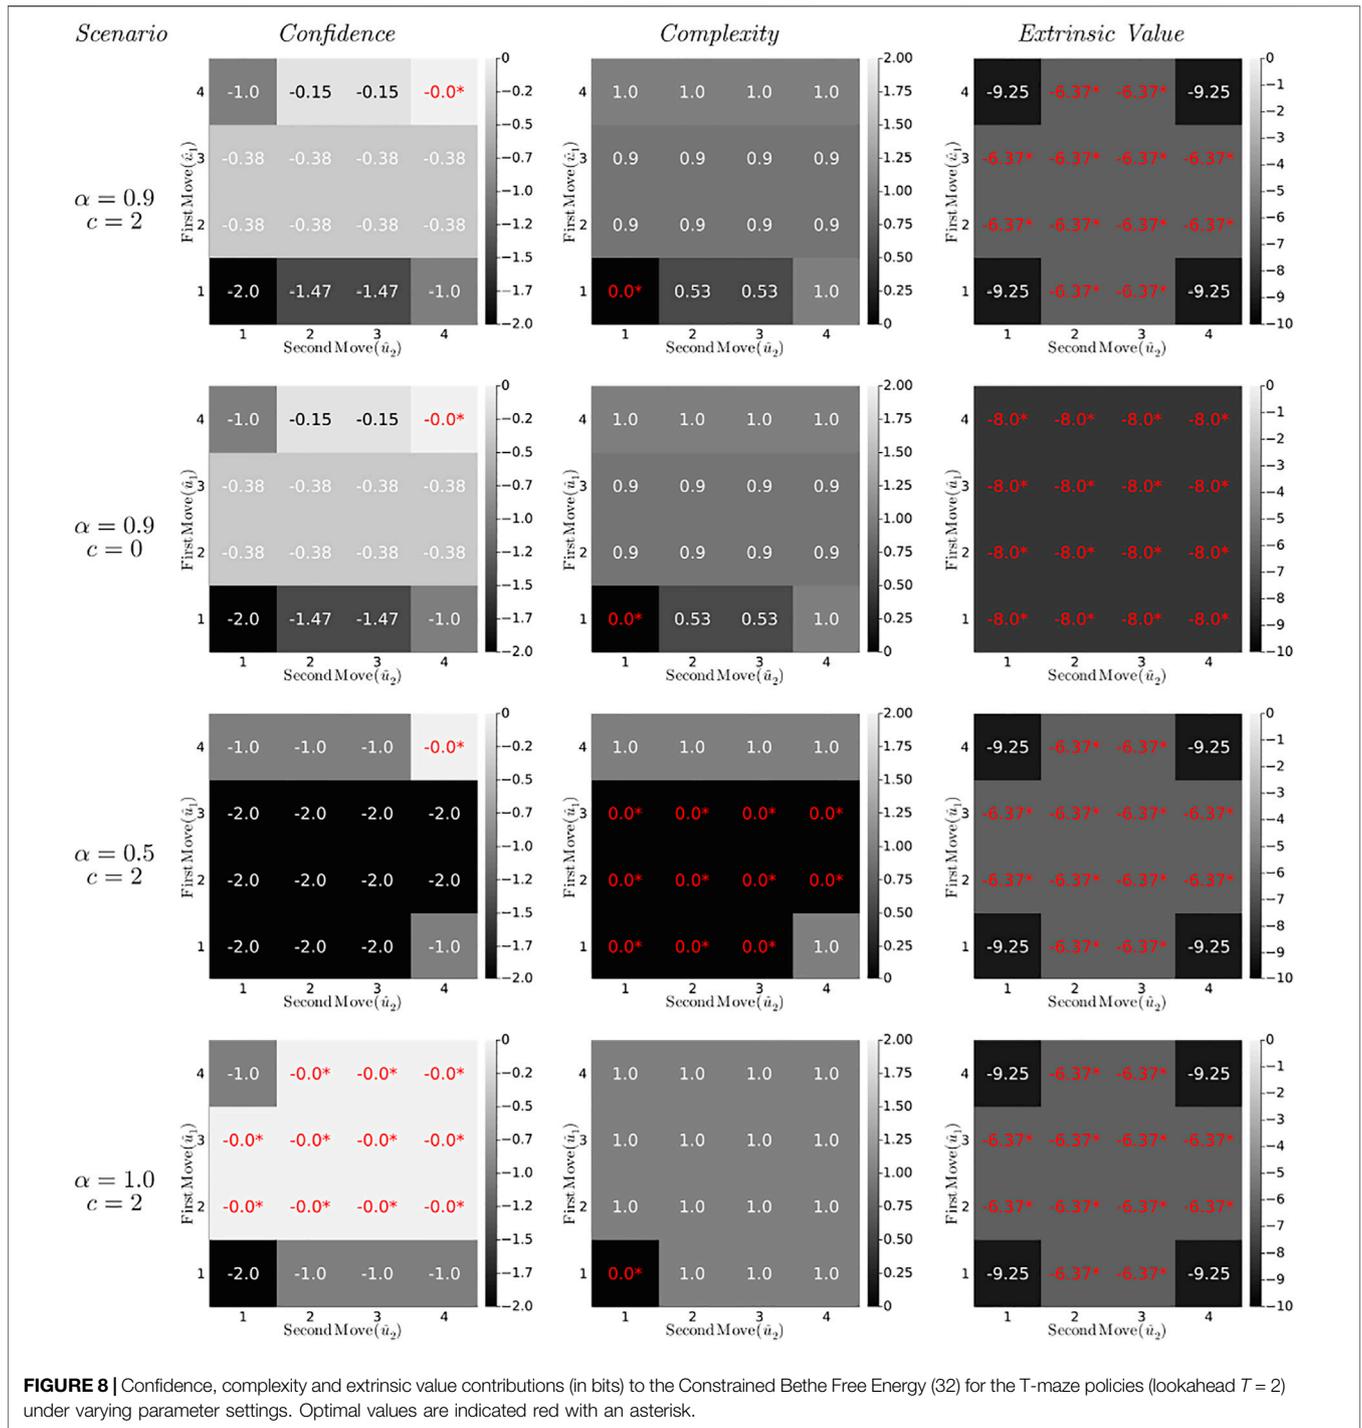
Simulated values for the CBFE decomposition (**Eq. 32**) in the T-maze application are shown in **Figure 8**, for four different T-maze scenarios. We summarize the most important observations below.

The first column of diagrams in **Figure 8** represents the confidence (**Eq. 34**) of the CBFE objective for all (planned) policies. The confidence prefers (or ties) the most informative policy (4,4) for all scenarios.

- In the first three scenarios (first three rows of diagrams), all policies other than (4,4) dismiss the opportunity to obtain full information about outcomes on two occasions ($T = 2$). This is reflected by a negative confidence value, which measures the average rejected information in bits. For example, the policy (1,1) rejects two possibilities to obtain 1 bit of information, leading to a confidence of -2 .

⁵ForneyLab is available at <https://github.com/biaslab/ForneyLab.jl>.

⁶Simulation source code is available at https://biaslab.github.io/materials/epistemic_search.zip.



- A change in the external value parameter c does not affect the confidence, which supports the interpretation of the confidence as an intrinsic quantity (Eq. 35).
- In the final scenario, the greedy policies (moving first to position 2 or 3) are on equal footing with the informative policies (moving first to 4). This is because in the final scenario, visiting position 2 or 3 offers the same amount of information (namely,

complete certainty) about the reward position, as would visiting the cue position.

The complexity (second column of diagrams) opposes changes in state beliefs that are unwarranted by the policy-induced state transitions, and guards against premature convergence of the state precision (Table 2). As a result, the complexity prefers (or ties) the most conservative policy (1,1) for all scenarios.

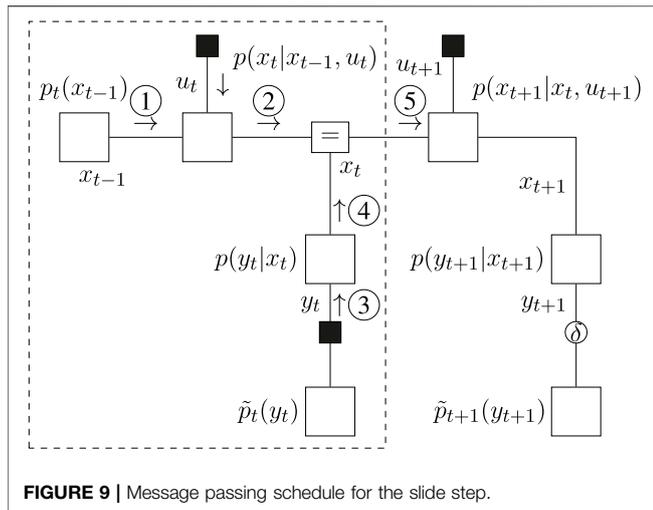


FIGURE 9 | Message passing schedule for the slide step.

- The complexity is unaffected by changes in utility (similar to the confidence), which supports the interpretation of the complexity as an intrinsic quantity (Eq. 35).
- In the third scenario ($\alpha = 0.5$), the greedy policies become tied in complexity with (most of) the ignorant policies. Because neither visiting a reward arm nor remaining at the initial position offers any useful information about the reward position, the state belief remains unaltered, and these policies incur no complexity penalty.

The extrinsic value (third column of diagrams) represents the value of external reward, and leads the agent to pursue extrinsically rewarding states.

- The extrinsic value is unaffected by changes in the epistemic reward probability parameter α , which supports the interpretation of the extrinsic value as an externally determined quantity.
- In the second scenario the reward utility vanishes ($c = 0$), and the extrinsic value becomes indifferent about policies.

7 INTERACTIVE SIMULATION

In this section we compare the resulting behavior of the CBFE agent with a traditional EFE agent, in *interaction* with a simulated environment.

7.1 Experimental Protocol

The experimental protocol governs how the agent interacts with its environment. In our protocol, the action and outcome at time t are the only quantities that are exchanged between the agent and the environment (generative process). The task of the agent is then to plan for actions that lead the agent to desired states. We adapt the experimental protocol of (van de Laar and de Vries, 2019) for the purpose of the current simulation. We write the model f_t with a time-subscript to indicate the time-dependent

statistics of the state prior as a result of the perceptual process (Section 2.2). The experimental protocol (Algorithm 1) then consists of five steps per time t .

Algorithm 1. Experimental Protocol

```

Given a model  $f_t$  with initial state and goal priors
for  $t = 1$  to  $N$  do
   $\hat{u}_t^* = \text{plan}(f_t)$  # Execute the planning algorithm
   $\hat{u}_t^* = \text{act}(\hat{u}_t^*)$  # Select the first action
  execute( $\hat{u}_t^*$ ) # Execute the action in the simulated environment
   $\hat{y}_t = \text{observe}()$  # Observe the new environmental outcome
   $f_{t+1} = \text{slide}(\hat{u}_t^*, \hat{y}_t)$  # Prepare the model for the next iteration
end for
    
```

The plan step solves the inference for planning (Section 2.3), and returns the active policy \hat{u}_t^* that represents the (believed) optimal sequence of future controls. In the act step, the first action \hat{u}_t^* is picked from the policy. The execute step then subsequently executes this action in the simulated environment. Execution will alter the state of the environment. In the observe step, the environment responds with a new observation \hat{y}_t . Given the action and resulting observation, the slide step then solves the inference for perception (Section 2.2) and prepares the model for the next step.

Inference for the slide step is illustrated in Figure 9, where message ③ propagates an observed outcome \hat{y}_t , and where message ⑤ summarizes the information contained within in the dashed box. Only the dashed sub-model is relevant to the slide step, that is, beliefs about the future do not influence ⑤. After computation, message ⑤ is normalized, and the resulting state posterior $q^*(x_t)$ is subsequently used as a prior to construct the model f_{t+1} for the next time-step, see also (van de Laar and de Vries, 2019).

7.2 Results for Interactive Simulation

We initialize an environment with the reward in the right arm (position 3). We then execute the experimental protocol of Algorithm 1, with lookahead $T = 2$, for $N = 2$ moves, on a dense landscape of varying reward probabilities α and utilities c (scenarios). After the first move, the environment returns an observations to the agent, which informs the agent about second move. After the second move, the expected reward that is associated with the resulting position is reported. We perform 10 simulations per scenario, and compute the average reward probability. The results of Figure 10 compare the average rewards of the CBFE agent and the EFE agent.

7.3 CBFE Agent EFE Agent

From the results of Figure 10 it can be seen that the region of zero average reward (dark region in lower left corner) is significantly smaller for the CBFE agent than for the EFE agent. This indicates that the CBFE agent accrues reward in a significantly larger portion of the scenario landscape than the EFE agent. In the lower left corner, the resulting CBFE agent trajectory becomes (4, 4), whereas the EFE agent trajectory becomes (4, 1). Although both agents observe the cue after their first move, they do not visit the indicated reward position in the second move, which leads to zero average reward. Note that neither objective is explicitly designed to optimize for average

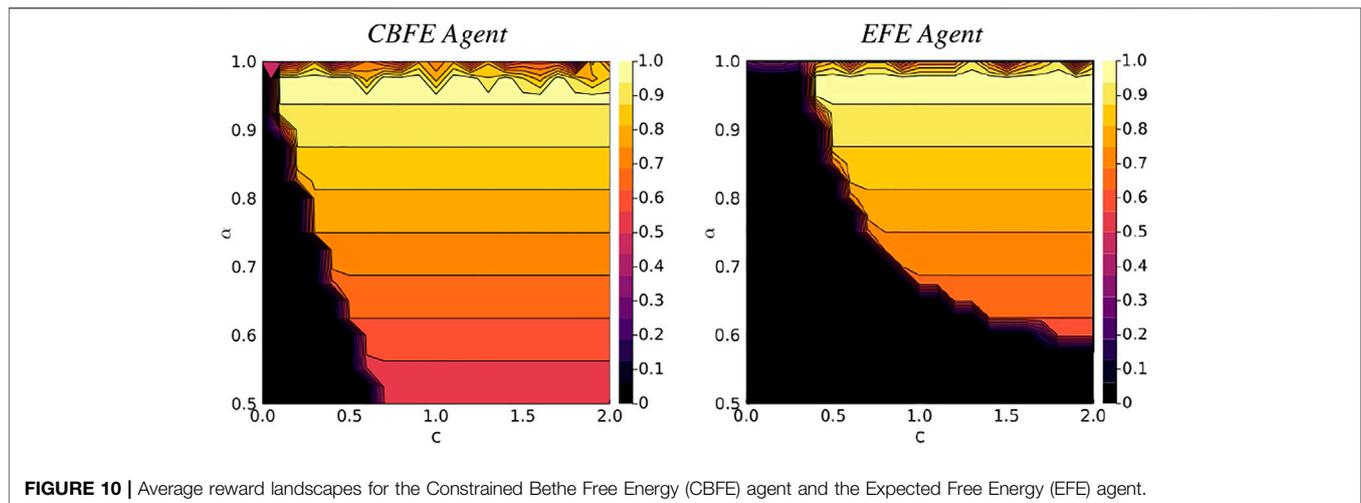


FIGURE 10 | Average reward landscapes for the Constrained Beth Free Energy (CBFE) agent and the Expected Free Energy (EFE) agent.

reward; both define a free energy instead, where multiple simultaneous forces are at play.

In the upper right regions, with high reward probability and utility, both agents consistently execute (4, 3). With this trajectory, the cue is observed after the first move, and the indicated (correct) reward position is visited in the second move, leading to an average reward of α . For reward probabilities close to $\alpha = 1$ however, the performance of both agents deteriorates. In this upper region, the informative policies become tied with the greedy policies (see **Figure 8**), and there is no single dominant trajectory. In some trajectories the agent enters the wrong arm on the first move, from which the agent cannot escape, and the average reward deteriorates.

Greedy behavior is also observed for the CBFE agent when informative priors c_k (**Eq. 42a**, **Eq. 42b**) are set for all k (including $k = 1$), conforming with the configuration of (Friston et al., 2015). With this configuration, expected reward for the CBFE agent deteriorates to 0.5 in the otherwise rewarding region. Interestingly, this change of priors does not affect results for the EFE agent. The resulting change in behavior suggests that the CBFE agent is more susceptible to temporal aspects of the goal prior configuration. While this effect may be considered a nuisance in some cases, it also allows for increased flexibility when assigning explicit temporal requirements about goals. For example, assigning an informative versus a flat prior for $k = 1$ respectively encodes an urgency in obtaining immediate reward versus a freedom to explore.

8 DISCUSSION

In this paper, we focused on epistemic drivers for behavior. We noted that the nature of the epistemic drive differs between an EFE and CBFE agent. Namely, the epistemic drive for the EFE agent stems directly from maximizing a mutual information term between states and observations (**Eq. 20**), while the epistemic drive for the CBFE agent stems from a self-evidencing mechanism (**Section 4.2**). In order to better understand the strengths and limitations of the driving forces for the CBFE, it

would be interesting to investigate its behavior in more challenging setups, including continuous variables, inference for control (van de Laar, 2018), and the effects point-mass constraints on other model variables.

Recent work by (Da Costa et al., 2020b; Millidge et al., 2021) shows that epistemic behavior does not occur when the goal prior goes to a point-mass. The work of (Millidge et al., 2021) points to the entropy of the observed variables $H[q(\mathbf{y})]$ as a pivotal quantity for epistemic behavior. The CBFE however does not include an entropy over observations, and still exhibits epistemic qualities. The difference in methods lies with the constraint quantity; namely (Da Costa et al., 2020a; Millidge et al., 2021) constrain the goal prior $\tilde{p}(\mathbf{y}) = \delta(\mathbf{y} - \hat{\mathbf{y}})$, while the current paper constrains the variational distribution $q(\mathbf{y}) = \delta(\mathbf{y} - \hat{\mathbf{y}})$ instead. While both constraints remove $H[q(\mathbf{y})]$ from the resulting FE objective, optimization of $\hat{\mathbf{y}}$ in the CBFE still induces an epistemic drive (**Section 4**). Our results thus show that epistemic drives for AIF prove to be more subtle than initially anticipated.

Our presented approach is uniquely scalable, because it employs off-the-shelf message passing algorithms. All message computations are local, which makes our approach naturally amenable to both parallel and on-line processing (Bagaev and de Vries, 2021). Especially AIF in deep hierarchical models might benefit from the improved computational properties of the CBFE. It will be interesting to investigate how the presented approach generalizes to more demanding (practical) settings.

As a generic variational inference procedure, the CBFE approach applies to arbitrary models. This allows researchers to investigate epistemics in a much wider class of models than previously available. One immediate avenue for further research is the integration of CBFE with predictive coding schemes (Friston and Kiebel, 2009; Bogacz, 2017; Millidge et al., 2020a). Predictive coding has so far been driven mainly by minimizing free energy in hierarchical models under the Laplace approximation. Here, the CBFE approach readily applies as well (Şenöz et al., 2021), allowing researchers to explore the effects of augmenting existing predictive coding models with epistemic components.

The derivation of alternative functionals that preserve the desirable epistemic behavior of EFE optimization is an active research area (Tschantz et al., 2020a; Sajid et al., 2021). There have been several interesting proposals such as the Free Energy of the Expected Future (Millidge et al., 2020b; Tschantz et al., 2020b; Hafner et al., 2020) or Generalized Free Energy (Parr and Friston, 2019), as well as amortization strategies (Ueltzhöffer, 2018; Millidge, 2019) and sophisticated schemes (Friston et al., 2021). Comparing behavior between the CBE and other free energy objectives might therefore prove an interesting avenue for future research.

In the original description of active inference, a policy precision is optimized during policy planning, and the policy for execution is sampled from a distribution of precision-weighted policies (Friston et al., 2015). The present paper does not consider precision optimization, and effectively assumes a large, fixed precision instead. In practice, this procedure consistently selects the policy with minimal free energy; see also *maximum selection* (in terms of value) as described by (Schwöbel et al., 2018). To accommodate for precision optimization, the CBE objective might be extended with a temperature parameter, mimicking thermodynamic descriptions of free energy (Ortega and Braun, 2013). Optimization of the temperature parameter might then relate to optimization of the policy precision, as often seen in biologically plausible formulations of AIF (FitzGerald et al., 2015).

Another interesting avenue for further research would be the design of a meta-agent that determines the statistics and temporal configuration of the goal priors. In our experiments we design the goal priors (Eq. 42a, Eq. 42b) ourselves, such that the agent is free to explore in the first move and seeks reward on the second move. The challenge then becomes to design a synthetic meta-agent that automatically generates an effective lower-level goal sequence from a single higher-level goal definition.

9 CONCLUSION

In this paper we presented mathematical arguments and simulations that show how inclusion of point-mass constraints on the Bethe Free Energy (BFE) leads to epistemic behavior. The thus obtained Constrained Bethe Free Energy (CBFE) has direct connections with formulations of the principle of least action in physics (Caticha, 2012), and can be conveniently optimized by message passing on a graphical representation of the generative model (GM).

Simulations for the T-maze task illustrate how a CBFE agent exhibits an epistemic drive, whereas the BFE agent lacks epistemic qualities. The key intuition behind the working mechanism of the CBFE is that point-mass constraints on observation variables explicitly encode the assumption that the agent will observe in the future. Although the actual value of these observation remains unknown, the agent “knows” that it will observe in the future, and it “knows” (through the GM) how these (potential) outcomes will influence inferences about states.

We dissected the CBFE objective in terms of its constituent drivers for behavior. In the CBFE framework, in addition to being functionals of the state beliefs, the confidence and complexity are viewed as functions of the potential outcomes and policy respectively. Simultaneous optimization of variational distributions and potential outcomes then leads the agent to prefer epistemic policies. Interactive simulations for the T-maze showed that, compared to an EFE agent, the CBFE agent incurs expected reward in a significantly larger portion of the scenario landscape.

We performed our simulations by message passing on a Forney-style factor graph representation of the generative model. The modularity of the graphical representation allows for flexible model search, and message passing allows for distributed computations that scale well to bigger models. Constraining the BFE and optimizing the CBFE objective by message passing thus suggests a simple and general mechanism for epistemic-aware AIF in free-form generative models.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/**Supplementary Material**, further inquiries can be directed to the corresponding author.

AUTHOR CONTRIBUTIONS

The original idea was conceived by TvdL. All authors contributed to further conceptual development of the methods that are presented in this manuscript. Simulations were performed by TvdL and MK. All authors contributed to writing the manuscript.

FUNDING

This research was made possible by funding from GN Hearing A/S. This work is part of the research programme Efficient Deep Learning with project number P16-25 project 5, which is (partly) financed by the Netherlands Organisation for Scientific Research (NWO).

ACKNOWLEDGMENTS

The authors gratefully acknowledge stimulating discussions with Dimitrije Marković of the Neuroimaging group at TU Dresden. The authors also thank the two reviewers for their valuable comments.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frobt.2022.794464/full#supplementary-material>

REFERENCES

- Bagaev, D., and de Vries, B. (2021). Reactive Message Passing for Scalable Bayesian Inference. arXiv:2112.13251 [cs] ArXiv: 2112.13251.
- Bogacz, R. (2017). A Tutorial on the Free-Energy Framework for Modelling Perception and Learning. *J. Math. Psychol.* 76, 198–211. doi:10.1016/j.jmp.2015.11.003
- Caticha, A. (2012). *Entropic Inference and the Foundations of Physics*. Sao Paulo, Brazil: EBEB-2012, the 11th Brazilian Meeting on Bayesian Statistics.
- Caticha, A. (2004). Relative Entropy and Inductive Inference. *AIP Conf. Proc.* 707, 75–96. ArXiv: physics/0311093. doi:10.1063/1.1751358
- Champion, T., Grzes, M., and Bowman, H. (2021). Realising Active Inference in Variational Message Passing: the Outcome-Blind Certainty Seeker. arXiv: 2104.11798 [cs] ArXiv: 2104.11798.
- Cox, M., van de Laar, T., and de Vries, B. (2019). A Factor Graph Approach to Automated Design of Bayesian Signal Processing Algorithms. *Int. J. Approximate Reasoning* 104, 185–204. doi:10.1016/j.ijar.2018.11.002
- Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., and Friston, K. (2020a). Active Inference on Discrete State-Spaces: a Synthesis. arXiv:2001.07203 [q-bio] ArXiv: 2001.07203. doi:10.1016/j.jmp.2020.102447
- Da Costa, L., Sajid, N., Parr, T., Friston, K., and Smith, R. (2020b). The Relationship between Dynamic Programming and Active Inference: the Discrete, Finite-Horizon Case. arXiv:2009.08111 [cs, math, q-bio] ArXiv: 2009.08111.
- Dauwels, J., Korl, S., and Loeliger, H.-A. (2005). “Expectation Maximization as Message Passing,” in *International Symposium on Information Theory. 2005. ISIT 2005. Proceedings*, Adelaide, SA (IEEE), 583–586. doi:10.1109/ISIT.2005.1523402
- Dauwels, J. (2007). “On Variational Message Passing on Factor Graphs,” in *IEEE International Symposium on Information Theory (Nice, France: IEEE)*, 2546–2550. doi:10.1109/ISIT.2007.4557602
- de Vries, B., and Friston, K. J. (2017). A Factor Graph Description of Deep Temporal Active Inference. *Front. Comput. Neurosci.* 11. doi:10.3389/fncom.2017.00095
- FitzGerald, T. H. B., Dolan, R. J., and Friston, K. (2015). Dopamine, Reward Learning, and Active Inference. *Front. Comput. Neurosci.* 9. doi:10.3389/fncom.2015.00136
- Forney, G. D. (2001). Codes on Graphs: normal Realizations. *IEEE Trans. Inform. Theor.* 47, 520–548. doi:10.1109/18.910573
- Friston, K., Da Costa, L., Hafner, D., Hesp, C., and Parr, T. (2021). Sophisticated Inference. *Neural Comput.* 33, 713–763. doi:10.1162/neco_a_01351
- Friston, K. J., Daunizeau, J., Kilner, J., and Kiebel, S. J. (2010). Action and Behavior: a Free-Energy Formulation. *Biol. Cybern.* 102, 227–260. doi:10.1007/s00422-010-0364-z
- Friston, K., and Kiebel, S. (2009). Predictive Coding under the Free-Energy Principle. *Phil. Trans. R. Soc. B* 364, 1211–1221. doi:10.1098/rstb.2008.0300
- Friston, K., Kilner, J., and Harrison, L. (2006). A Free Energy Principle for the Brain. *J. Physiology-Paris* 100, 70–87. doi:10.1016/j.jphysparis.2006.10.001
- Friston, K., and Penny, W. (2011). Post Hoc Bayesian Model Selection. *Neuroimage* 56, 2089–2099. doi:10.1016/j.neuroimage.2011.03.062
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., and Pezzulo, G. (2015). Active Inference and Epistemic Value. *Cogn. Neurosci.* 6, 187–214. doi:10.1080/17588928.2015.1020053
- Friston, K., Schwartenbeck, P., Fitzgerald, T., Moutoussis, M., Behrens, T., and Dolan, R. J. (2013). The Anatomy of Choice: Active Inference and agency. *Front. Hum. Neurosci.* 7, 598. doi:10.3389/fnhum.2013.00598
- Hafner, D., Ortega, P. A., Ba, J., Parr, T., Friston, K., and Heess, N. (2020). Action and Perception as Divergence Minimization. arXiv:2009.01791 [cs, math, stat] ArXiv: 2009.01791.
- Hohwy, J. (2021). “Conscious Self-Evidencing,” in *Review of Philosophy and Psychology* (Publisher: Springer), 1–20. doi:10.1007/s13164-021-00578-x
- Hohwy, J. (2016). The Self-Evidencing Brain. *Nous* 50, 259–285. doi:10.1111/nous.12062
- Koller, D., and Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, Massachusetts: MIT press.
- Korl, S. (2005). *A Factor Graph Approach to Signal Modelling*. system identification and filtering. Zurich: Swiss Federal Institute of Technology. Ph.D. thesis.
- Lanczos, C. (2012). *The Variational Principles of Mechanics*. Chelmsford, Massachusetts: Courier Corporation.
- Loeliger, H.-A., Dauwels, J., Hu, J., Korl, S., Ping, L., and Kschischang, F. R. (2007). The Factor Graph Approach to Model-Based Signal Processing. *Proc. IEEE* 95, 1295–1322. doi:10.1109/JPROC.2007.896497
- Loeliger, H.-A. (2002). “Least Squares and Kalman Filtering on Forney Graphs,” in *Codes, Graphs, and Systems*. Editors R. E. Blahut and R. Koetter (Boston, MA: Springer US), 670, 113–135. doi:10.1007/978-1-4615-0895-3_7
- Loeliger, H. (2004). An Introduction to Factor Graphs. *IEEE Signal. Process. Mag.* 21, 28–41. doi:10.1109/MSP.2004.1267047
- Millidge, B. (2019). Deep Active Inference as Variational Policy Gradients. arXiv: 1907.03876 [cs] ArXiv: 1907.03876
- Millidge, B., Tschantz, A., and Buckley, C. L. (2020a). Predictive Coding Approximates Backprop along Arbitrary Computation Graphs. arXiv preprint arXiv:2006.04182
- Millidge, B., Tschantz, A., and Buckley, C. L. (2020b). Whence the Expected Free Energy? arXiv preprint arXiv:2004.08128.
- Millidge, B., Tschantz, A., Seth, A., and Buckley, C. (2021). Understanding the Origin of Information-Seeking Exploration in Probabilistic Objectives for Control. arXiv preprint arXiv:2103.06859.
- Ortega, P. A., and Braun, D. A. (2013). Thermodynamics as a Theory of Decision-Making with Information-Processing Costs. *Proc. R. Soc. A.* 469, 20120683. doi:10.1098/rspa.2012.0683
- Parr, T., and Friston, K. J. (2019). Generalised Free Energy and Active Inference. *Biol. Cybern.* 113, 495–513. Publisher: Springer. doi:10.1007/s00422-019-00805-w
- Parr, T., Markovic, D., Kiebel, S. J., and Friston, K. J. (2019). Neuronal Message Passing Using Mean-Field, Bethe, and Marginal Approximations. *Sci. Rep.* 9, 1889. doi:10.1038/s41598-018-38246-3
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Sajid, N., Faccio, F., Da Costa, L., Parr, T., Schmidhuber, J., and Friston, K. (2021). Bayesian Brains and the Renyi Divergence. arXiv preprint arXiv: 2107.05438.
- Schwöbel, S., Kiebel, S., and Marković, D. (2018). Active Inference, Belief Propagation, and the Bethe Approximation. *Neural Comput.* 30, 2530–2567. doi:10.1162/neco_a_01108
- Şenöz, İ., van de Laar, T., Bagaev, D., and de Vries, B. d. (2021). Variational Message Passing and Local Constraint Manipulation in Factor Graphs. *Entropy* 23, 807. Publisher: Multidisciplinary Digital Publishing Institute. doi:10.3390/e23070807
- Tschantz, A., Baltieri, M., Seth, A. K., and Buckley, C. L. (2020a). “Scaling Active Inference,” in 2020 International Joint Conference on Neural Networks (IJCNN) (IEEE), 1–8. doi:10.1109/ijcnn48605.2020.9207382
- Tschantz, A., Millidge, B., Seth, A. K., and Buckley, C. L. (2020b). Reinforcement Learning through Active Inference. arXiv:2002.12636 [cs, eess, math, stat] ArXiv: 2002.12636.
- Ueltzhöffer, K. (2018). Deep Active Inference. *Biol. Cybern.* 112 (6), 547–573. doi:10.1007/s00422-018-0785-7
- van de Laar, T. (2019). *Automated Design of Bayesian Signal Processing Algorithms*. Eindhoven, Netherlands: Eindhoven University of Technology. Ph.D. thesis.
- van de Laar, T., Şenöz, İ., Özçelikkale, A., and Wymeersch, H. (2021). Chance-Constrained Active Inference. arXiv preprint arXiv:2102.08792. doi:10.1162/neco_a_01427
- [Dataset] van de Laar, T. (2018). Simulating Active Inference Processes with Message Passing. *Front. Robot. AI* 6. doi:10.3389/frobt.2019.00020
- van de Laar, T. W., and de Vries, B. (2019). Simulating Active Inference Processes by Message Passing. *Front. Robot. AI* 6, 20. doi:10.3389/frobt.2019.00020
- Wainwright, M. J., and Jordan, M. I. (2007). Graphical Models, Exponential Families, and Variational Inference. *FNT Machine Learn.* 1, 1–305. doi:10.1561/22000000001
- Yedidia, J. S., Freeman, W. T., and Weiss, Y. (2005). Constructing Free-Energy Approximations and Generalized Belief Propagation Algorithms. *IEEE Trans. Inform. Theor.* 51, 2282–2312. doi:10.1109/TIT.2005.850085
- Zhang, D., Wang, W., Fettweis, G., and Gao, X. (2017). Unifying Message Passing Algorithms under the Framework of Constrained Bethe Free Energy Minimization. arXiv:1703.10932 [cs, math] ArXiv: 1703.10932.

Conflict of Interest: BdV was employed by the GN Hearing Benelux BV. MK was employed by Nested Minds Network Ltd.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in

this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 van de Laar, Koudahl, van Erp and de Vries. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.