

INTERNATIONAL WORKSHOP ON ACTIVE INFERENCE · 2026


Expected free energy as an information constraint on the Bethe Lagrangian

Wouter M. Kouw

Bayesian Intelligent Autonomous Systems lab
Electrical Engineering | Eindhoven University of Technology

Code: github.com/biaslab/IWAI2026-EFEasBetheConstraint

Navigate to an uncertain goal



The agent must move from its start to the goal.

The goal is left (L) or right (R),
but which one is uncertain.


The cue contains information to go left or right.

PROBLEM SETTING

Variables and the planning model

Generative model for planning, with a goal prior on every predicted observation

$$p(y_\tau, u_\tau, x_\tau, x_t, \Theta | \mathcal{D}_t) \propto \underbrace{p(x_t | \mathcal{D}_t)}_{f_t} \underbrace{p(B | \mathcal{D}_t)}_{f_B} \underbrace{p(A | \mathcal{D}_t)}_{f_A} \prod_{k \in \tau} \underbrace{p(x_k | x_{k-1}, u_k, B)}_{f_x} \underbrace{p(y_k | x_k, A)}_{f_y} \underbrace{p(u_k)}_{f_u} \underbrace{p^*(y_k)}_{f_*}$$



$x_k \in \mathcal{X}$

hidden state

$u_k \in \mathcal{U}$

action

$y_k \in \mathcal{Y}$

observation

$\Theta = \{A, B\}$

model parameters

$u_\tau = (u_{t+1}, \dots, u_{t+K})$

policy

$p^*(y_k)$

goal prior

$\mathcal{D}_t = \{y_{1:t}, u_{1:t}\}$

data so far

THE MISMATCH

Bethe free energy lacks the information gain

Bethe free energy

$$B[q] = \underbrace{\sum_{a \in \mathcal{V}} \mathbb{E}_{q_a} \left[\ln \frac{q_a(s_a)}{f_a(s_a)} \right]}_{\text{factor nodes}} + \underbrace{\sum_{i \in \mathcal{E}} \mathbb{H}[q_i(s_i)]}_{\text{edges}}$$

Applying form constraints simplifies it

$$B[q] = \mathbb{D}_{\text{KL}}[q(u_\tau) \parallel p(u_\tau)] - \sum_{k \in \tau} \mathbb{E}_q [\ln p^*(y_k)]$$

$$q_{\text{EFE}}(u_\tau) \propto p(u_\tau) \exp \left(\sum_{k \in \tau} \underbrace{\mathbb{E}_{q(y_k | u_\tau)} [\ln p^*(y_k)]}_{\text{goal-seeking}} + \underbrace{\mathbb{I}[x_k, \Theta; y_k | u_\tau]}_{\text{information-seeking}} \right)$$

$$q_{\text{Bethe}}(u_\tau) \propto p(u_\tau) \exp \left(\sum_{k \in \tau} \mathbb{E}_{q(y_k | u_\tau)} [\ln p^*(y_k)] \underbrace{\left[+ \mathbb{I}[x_k, \Theta; y_k | u_\tau] \right]}_{\text{missing}} \right)$$

Idea: constrain the Bethe functional so that information gain re-appears.

DEFINITION · FOR EVERY PLANNING STEP $k \in T$

The information constraint

$$\underbrace{\mathbb{E}_{q(u_\tau)} \left[\mathbb{I}[x_k, \Theta; y_k | u_\tau] \right]}_{\text{expected information gain}} \geq \underbrace{\mathbb{H}[p^*(y_k)]}_{\text{floor } \beta_k}$$

Floor from the goal prior

$$\beta_k = \mathbb{H}[p^*(y_k)]$$

- No new hyperparameter.
- Broader preferences demand more information.

KKT multiplier γ

$$\gamma_k \geq 0, \quad \gamma_k \left(\beta_k - \mathbb{E}_{q(u_\tau)} [\mathbb{I}_k] \right) = 0$$

- Zero when the floor is met.
- The drive switches on only when needed.

THEOREM 1

From Lagrangian to stationary policy

Constrained Lagrangian: Bethe free energy, normalisation, and the information constraint in KKT form

$$L[q] = B[q] + \lambda \left(\sum_{u_\tau} q(u_\tau) - 1 \right) + \sum_{k \in \tau} \gamma_k \left(\beta_k - \mathbb{E}_{q(u_\tau)} [\mathbb{I}[x_k, \Theta; y_k | u_\tau]] \right)$$

Setting $\delta L / \delta q = 0$ and normalising: every stationary point has policy belief

$$q^*(u_\tau) \propto p(u_\tau) \exp \left(\underbrace{\sum_{k \in \tau} \mathbb{E}_{q(y_k | u_\tau)} [\ln p^*(y_k)]}_{\text{goal-seeking}} + \gamma_k \underbrace{\mathbb{I}[x_k, \Theta; y_k | u_\tau]}_{\text{information-seeking}} \right)$$


PROPOSITION 1 · EQUIVALENCE TO EFE

$$\gamma_k = 1 \quad \forall k \in \tau \quad \implies \quad q^*(u_\tau) \propto p(u_\tau) \exp(-G(u_\tau))$$


Information seeking switches off at the cue

t = 0


Goal unknown: heading for the cue




(a) Predicted goal location (mean $\pm$ 1 sd)




(b) Solved multiplier γ



(c) Leverage of each term (spread across the 729 policies, nats)



One sweep of messages per policy




① $q(x_k | u_\tau)$
Forward: predicted state trajectories.

② $q(y_k | u_\tau), \mathbb{I}_k$
Forward: predictions and information.

③ $p^*(y_k), \mathbb{E}[\ln p^*(y_k)]$
Backward: goal priors and the goal term.

④ $q^*(u_\tau)$
Backward: collected into the policy.

The constraint sends the agent to the cue



SUMMARY

Take-aways

Constrained Bethe

One can impose a constraint $\text{information gain} \geq \text{goal-prior entropy}$, that will re-introduces information-seeking terms into the posterior.

EFE exactly at $\gamma = 1$

The family of stationary policies contains the EFE policy as a special case.

γ solved, not tuned

The effect of the information-seeking term can be automatically turned off when actions under the goal-seeking term produce "enough" information.

DISCUSSION

Nesting: who sets the information demand?

A higher-level agent could operate this agent's exploration: by setting its KKT multiplier γ directly, or by setting the variance of its goal prior, and with it the information floor β . Exploration then becomes an action of the agent above.

Paper and code: github.com/biaslab/IWAI2026-EFEasBetheConstraint