

PlanningEFEMix: Hybrid Active Inference for Sequential Decision-Making under Uncertainty

Bhagyeshkumar Chokhawala¹, Atif Farid Mohammad²

¹Capitol Technology University, Laurel, Maryland, USA, bchokhawala@captechu.edu

²PhD, Capitol Technology University, Laurel, Maryland, USA, afmohammad@captechu.edu

Abstract: Sequential decision-making under uncertainty remains a key challenge in artificial intelligence, especially in environments marked by partial observability and noisy feedback. While reinforcement learning and probabilistic planning have achieved significant success, each approach has limitations when used alone, including instability under noisy conditions and reliance on accurate generative models. Active Inference offers a principled alternative by framing perception, learning, and action selection as the minimization of Expected Free Energy (EFE), unifying exploration and goal-directed behavior within a Bayesian framework (Friston, 2010; Friston et al., 2020). This paper introduces PlanningEFEMix, a hybrid decision-making algorithm that enables meta-level planning across diverse inference agents using Expected Free Energy as a shared objective. The framework combines deterministic Active Inference, POMDP-based belief updating, contrastive learning, and model-free reinforcement learning within a single planning loop. Candidate actions are assessed via forward simulation across agents and selected via a softmax policy augmented with an adaptive, state-dependent bias memory that incorporates experiential feedback. An experimental evaluation on a noisy preference inference benchmark shows improved robustness and stability compared to single-agent baselines, confirming the effectiveness of hybrid Active Inference planning under uncertainty.

Keywords: Active Inference, Expected Free Energy, Hybrid Decision-Making, Reinforcement Learning, POMDP, Meta-Inference, Sequential Planning, Uncertainty Handling

Introduction

Sequential decision-making under uncertainty is fundamental to intelligent behavior and underlies applications ranging from robotics and recommendation systems to adaptive control and human–AI interaction. Classical approaches typically fall into two categories: model-free reinforcement learning (RL), which optimizes policies through trial-and-error interaction (Sutton & Barto, 2018), and model-based planning approaches, such as Partially Observable Markov Decision Processes (POMDPs), which explicitly represent latent states and observation uncertainty (Kaelbling et al., 1998). Although powerful, both paradigms suffer from well-known limitations. Reinforcement learning often requires extensive data and can exhibit unstable behavior under noise, while POMDP-based approaches rely on accurate transition and observation models that are difficult to specify in complex environments.

Active Inference provides an alternative formulation of decision-making grounded in the Free Energy Principle, which posits that adaptive systems act to minimize variational free energy in order to maintain viable states (Friston, 2010). In this framework, action selection emerges from minimizing Expected Free Energy, a quantity that naturally balances epistemic uncertainty reduction and pragmatic goal satisfaction (Friston et al., 2020; Millidge et al., 2020). This unification of exploration and exploitation distinguishes Active Inference from classical reinforcement learning, where exploration is often introduced heuristically.

Despite its conceptual appeal, many practical implementations of Active Inference rely on a single generative model or decision strategy, limiting adaptability when model assumptions are violated. Empirical and theoretical work in neuroscience and machine learning suggests that intelligent agents often rely on multiple competing decision systems that are dynamically recruited based on context (Daw et al., 2005; Dayan & Daw, 2008). This motivates hybrid approaches that integrate multiple inference strategies while retaining a coherent decision objective.

This paper introduces **PlanningEFEMix**, a hybrid Active Inference planning algorithm that performs meta-level decision-making across heterogeneous inference agents. Rather than committing to a single model, the framework treats multiple agents as competing internal hypotheses and selects actions through Expected Free Energy minimization augmented by experience-driven bias learning.

Background

Active Inference and Expected Free Energy

Active Inference formalizes perception, learning, and action within a unified Bayesian framework derived from the Free Energy Principle (Friston, 2010). In discrete state spaces, agents maintain beliefs over latent states and select actions that minimize Expected Free Energy (Friston et al., 2020). Expected Free Energy can be expressed as:

$$G(a) = \mathbb{E}_{q(o,s|a)}[\ln q(s|a) - \ln p(o,s)]$$

This formulation decomposes into an epistemic term, capturing expected information gain, and a pragmatic term, reflecting alignment with preferred outcomes (Millidge et al., 2020). As a result, agents are incentivized both to reduce uncertainty and to pursue goal-consistent states without relying on explicit reward signals.

Active Inference has been successfully applied to planning and navigation tasks, demonstrating how forward simulation and policy evaluation emerge naturally from EFE minimization (Kaplan & Friston, 2018). However, these approaches typically assume a fixed generative model, limiting flexibility in complex or non-stationary environments.

Limitations of Single-Paradigm Decision Models

Model-free reinforcement learning has achieved remarkable success across many domains but is sensitive to observation noise and often requires extensive interaction data (Sutton & Barto, 2018). Model-based approaches, such as POMDPs, offer a principled handling of uncertainty but scale poorly and depend on accurate generative assumptions (Kaelbling et al., 1998). Deterministic Active Inference agents provide transparent planning mechanisms but lack mechanisms for adapting decision strategies based on long-term empirical performance (Sajid et al., 2021).

Comparative analyses have highlighted that no single decision paradigm dominates across all environments, motivating hybrid approaches that combine planning and learning mechanisms (Paul et al., 2024). PlanningEFEMix builds on this insight by integrating multiple inference agents within a unified Active Inference framework.

PlanningEFEMix Framework

Overview

PlanningEFEMix is a hybrid decision-making algorithm that performs meta-level planning across a set of heterogeneous inference agents:

$$\mathcal{A} = \{A_{DPEFE}, A_{POMDP}, A_{CL}, A_{RL}\}$$

Each agent embodies distinct inductive biases and decision mechanisms. Deterministic Active Inference agents emphasize information-seeking and goal alignment, POMDP-based agents focus on probabilistic belief updating under partial observability (Kaelbling et al., 1998), contrastive learning agents contribute representation-based discrimination capabilities (Chen et al., 2020), and reinforcement learning agents provide experience-driven value learning (Sutton & Barto, 2018).

State Representation and Bias Memory

Observations are encoded into abstract state representations s_t . To enable experience-driven adaptation, PlanningEFEMix maintains a state–action bias memory $B(s, a)$, updated according to:

$$B_{t+1}(s, a) = \lambda B_t(s, a) + (1 - \lambda)(-G_t(a))$$

This bias mechanism is inspired by value-updating processes observed in biological decision systems (Dayan & Daw, 2008) and enables the framework to favor actions that have historically reduced Expected Free Energy in similar contexts.

Meta-Level Planning and Action Selection

For each candidate action, Expected Free Energy estimates are computed independently by each agent and aggregated:

$$\tilde{G}(a) = \sum_i G_i(a) + B(s, a)$$

Action selection follows a softmax policy:

$$P(a | s) = \frac{\exp(-\tilde{G}(a)/\tau)}{\sum_{a'} \exp(-\tilde{G}(a')/\tau)}$$

This stochastic selection mechanism supports structured exploration while preventing premature convergence to suboptimal policies, a limitation commonly observed in greedy reinforcement learning strategies (Daw et al., 2005).

Algorithm Description

At each decision step, PlanningEFEMix encodes observations into an abstract state, simulates candidate actions across all inference agents, computes Expected Free Energy estimates, aggregates these values with state-dependent bias, and samples an action via a softmax policy. Following action execution, the bias memory is updated based on observed outcomes, enabling continual adaptation without explicit reward shaping.

This design operationalizes recent theoretical work on predictive planning and counterfactual learning in Active Inference, where agents dynamically switch between strategies based on expected informational and pragmatic value (Paul et al., 2024).

Experimental Evaluation

Setup

The framework is evaluated on a preference inference task under partial observability and varying levels of observation noise. Baseline agents include deterministic Active Inference, POMDP-based belief agents, contrastive learning agents, and model-free reinforcement learning agents. Similar benchmark settings have been used in prior comparative studies of decision-making architectures (Sajid et al., 2021).

Results

Across all noise regimes, PlanningEFEMix demonstrates superior robustness and stability compared to single-agent baselines. Performance degradation under increasing noise is gradual, reflecting the benefits of agent diversity and bias-driven adaptation. These findings are consistent

with prior evidence that hybrid decision systems outperform monolithic architectures in uncertain environments (Daw et al., 2005; Paul et al., 2024).

Ablation Studies

Ablation experiments reveal that removing bias memory or meta-level agent integration significantly degrades performance, particularly under high noise. Greedy action selection further reduces robustness, underscoring the importance of stochastic planning in Active Inference-based systems (Kaplan & Friston, 2018).

Discussion

PlanningEFEMix demonstrates that Active Inference can be operationalized as a hybrid planning framework that integrates multiple inference strategies while maintaining a unified decision objective. By combining model-based planning, representation learning, and model-free adaptation, the framework achieves robustness that neither paradigm provides individually. This aligns with growing evidence that intelligent systems benefit from combining multiple decision mechanisms rather than relying on a single optimal strategy (Pezzulo et al., 2024).

Limitations and Future Work

The primary limitation of PlanningEFEMix is the computational overhead associated with maintaining multiple inference agents. Prior work on deep Active Inference and Monte-Carlo planning suggests promising directions for scaling such frameworks through approximation and hierarchical modeling (Ueltzhöffer, 2018; Fountas et al., 2020). Future research will explore these avenues and extend the framework to continuous state and action spaces.

Conclusion

This paper presented PlanningEFEMix, a hybrid Active Inference planning algorithm that performs meta-level decision-making across heterogeneous inference agents using Expected Free Energy minimization and adaptive bias learning. Experimental results demonstrate improved robustness under uncertainty compared to single-agent baselines, supporting the viability of hybrid Active Inference systems for complex sequential decision-making.

References

- Botvinick, M., Niv, Y., & Barto, A. (2009). Hierarchically organized behavior and its neural foundations. *Cognition*, *113*(3), 262–280. <https://doi.org/10.1016/j.cognition.2008.08.011>
- Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 1597–1607.
- Da Costa, L., Lanillos, P., Sajid, N., Friston, K., & Khan, S. (2022). How active inference could help revolutionise robotics. *Entropy*, *24*(3), 361. <https://doi.org/10.3390/e24030361>
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704–1711. <https://doi.org/10.1038/nrn1560>
- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective & Behavioral Neuroscience*, *8*(4), 429–453. <https://doi.org/10.3758/CABN.8.4.429>
- Fountas, Z., Sajid, N., Mediano, P. A. M., & Friston, K. (2020). Deep active inference agents using Monte-Carlo methods. *arXiv preprint arXiv:2006.04176*.
- Friston, K. (2010). The free-energy principle: A unified brain theory. *Nature Reviews Neuroscience*, *11*(2), 127–138. <https://doi.org/10.1038/nrn2787>
- Friston, K., Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., & Hager, J. (2020). Active inference on discrete state-spaces: A synthesis. *Journal of Mathematical Psychology*, *99*, 102447. <https://doi.org/10.1016/j.jmp.2020.102447>
- Friston, K. J., Parr, T., & de Vries, B. (2017). The graphical brain: Belief propagation and active inference. *Network Neuroscience*, *1*(4), 381–414. https://doi.org/10.1162/NETN_a_00018

- Gershman, S. J. (2018). The successor representation: Its computational logic and neural substrates. *Journal of Neuroscience*, 38(33), 7193–7200. <https://doi.org/10.1523/JNEUROSCI.0151-18.2018>
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2), 99–134. [https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X)
- Kaplan, R., & Friston, K. (2018). Planning and navigation as active inference. *Biological Cybernetics*, 112(4), 323–343. <https://doi.org/10.1007/s00422-018-0753-2>
- Millidge, B., Tschantz, A., & Buckley, C. L. (2020). Whence the expected free energy? *arXiv preprint arXiv:2004.08128*.
- Parr, T., Pezzulo, G., & Friston, K. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press.
- Paul, A., Isomura, T., & Razi, A. (2024). On predictive planning and counterfactual learning in active inference. *Entropy*, 26(6), 484. <https://doi.org/10.3390/e26060484>
- Pezzulo, G., Parr, T., Cisek, P., Clark, A., & Friston, K. (2024). Generating meaning: Active inference and the scope and limits of passive AI. *Trends in Cognitive Sciences*, 28(2), 97–112. <https://doi.org/10.1016/j.tics.2023.10.002>
- Sajid, N., Ball, P. J., Parr, T., & Friston, K. (2021). Active inference: Demystified and compared. *Neural Computation*, 33(3), 674–712. https://doi.org/10.1162/neco_a_01375
- Silver, D., Sutton, R., & Müller, M. (2012). Reinforcement learning of local shape in the game of Go. *Proceedings of the AAAI Conference on Artificial Intelligence*, 26(1), 1053–1059.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
- Ueltzhöffer, K. (2018). Deep active inference. *Biological Cybernetics*, 112(6), 547–573. <https://doi.org/10.1007/s00422-018-0785-7>