

Active Inference for Autonomous Decision-Making with Contextual Multi-Armed Bandits

Shohei Wakayama and Nisar Ahmed*

Abstract—In autonomous robotic decision-making under uncertainty, the tradeoff between exploitation and exploration of available options must be considered. If secondary information associated with options can be utilized, such decision-making problems can often be formulated as contextual multi-armed bandits (CMABs). In this study, we apply active inference, which has been actively studied in the field of neuroscience in recent years, as an alternative action selection strategy for CMABs. Unlike conventional action selection strategies, it is possible to rigorously evaluate the uncertainty of each option when calculating the expected free energy (EFE) associated with the decision agent’s probabilistic model, as derived from the free-energy principle. We specifically address the case where a categorical observation likelihood function is used, such that EFE values are analytically intractable. We introduce new approximation methods for computing the EFE based on variational and Laplace approximations. Extensive simulation study results demonstrate that, compared to other strategies, active inference generally requires far fewer iterations to identify optimal options and generally achieves superior cumulative regret, for relatively low extra computational cost.

I. INTRODUCTION

In robotics, autonomous robots must often optimize choices from among multiple alternatives to accomplish a task under uncertainty. Such decision-making can be complex when the outcomes obtained for different options are stochastic and their distributions are unknown *a priori*. If a robot could devote unlimited time and resources to the decision-making process, it would be possible to simulate all options a significant number of times and derive accurate outcome distributions before finalizing the best option. However, in many environments where autonomous robots are actually expected to be deployed, time and resource (e.g. bandwidth and energy) constraints are severe. This can also hamper the ability of human users to effectively interact with robots. In such settings, lightweight and effective algorithms are needed for decision-making with limited information.

For instance, as illustrated in Fig. 1, suppose a robot in a remote environment (such as an icy solar system moon) is assigned a task of collecting scientifically interesting objects (e.g. ice deposits and minerals) from a single search area using a resource-intensive manipulator. Although the robot could send image and spectral data about each area to human experts and let them decide where to investigate, mission time and communication constraints (e.g. due to

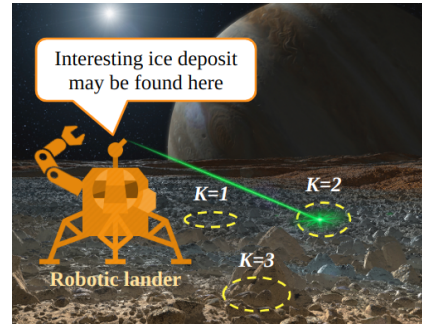


Fig. 1. A robotic lander reasons about whether candidate dig sites on an icy moon contain scientifically interesting data. To speed up the process, it is desirable not only to balance exploitation and exploration, but also incorporate a selection bias, i.e. prior preferences, for desired observations.

high radiation and hours long round trip time delays for speed of light signals) make smooth human-robot interaction hard in such missions [1], [2]. Thus, the robot must autonomously and accurately infer the outcome probability that each area contains valuable objects, while using sensors that are easily deployable and consume little energy. To improve efficiency, the robot must strike a balance between increasing the certainty of plausible areas (*exploitation*) and reducing the uncertainty of unknown areas (*exploration*), taking into account the surrounding environment, mission status, etc.

Such a decision-making problem can be formulated as a contextual multi-bandit (CMAB), a type of reinforcement learning problem which has originally been studied and applied in recommendation systems [3] and recently utilized in robotics [4], [5]. Generally speaking, however, bandit problems (depending on their scale and complexity) typically require a large number of iterative interactions with the problem environment to finalize an optimal option, which can be a bottleneck in applying them to many practical robotic applications like space exploration. Therefore, to address such concerns, this paper considers *active inference* [6], [7] as an action selection strategy for CMAB problems. Active inference has recently attracted attention in the field of neuroscience because it can mathematically rigorously take into account the uncertainty of each option in the process by calculating a measure named *expected free energy (EFE)* and incorporate a selection bias toward preferred outcomes. In this way, it is theoretically possible to determine the best option with fewer iterations. In fact, active inference has been applied to standard stationary and dynamic switching MABs and those properties have been experimentally verified [8], but no theory or studies have yet applied it to the important case of contextual bandits. Also, for CMABs with discrete

*Work supported by the NASA COLDTech Program, grant #80NSSC21K1031. S. Wakayama was also supported by the Masason Foundation. The authors are with the Smead Aerospace Engineering Sciences Department, University of Colorado Boulder, Boulder, CO 80303 USA [shohei.wakayama; nisar.ahmed}@colorado.edu

categorical action outcomes, the EFE terms are generally not tractable for general context-dependent likelihood models. Thus, we propose two novel statistical EFE approximation methods based on variational inference and Laplace approximation, to permit the use of active inference for solving CMABs. These methods are demonstrated and validated in a simulated autonomous search area selection problem, and their performances are compared with existing state of the art CMAB action selection strategies. For the remainder of the paper: Sec. II reviews the MAB and CMAB problems, as well as an overview of active inference and the free-energy principle behind the theory; Sec. III formulates the action selection strategy for CMAB problems using active inference, and derives the EFE approximations; Sec. IV presents the setup and results of the simulation experiment (motivated by the autonomous remote exploration example); and Sec. V presents conclusions.

II. BACKGROUND AND RELATED WORK

A. Multi Armed Bandits (MABs) and Contextual MABs

The multi-armed bandit (MAB) is a classic reinforcement learning problem of optimal selection from multiple alternatives with uncertain outcomes. The MAB allows a decision-making agent like an autonomous robot to account for the tradeoff between exploitation (i.e. preferring alternatives with the higher expected outcome) and exploration (i.e. preferring alternatives with a more uncertain distribution of outcomes) [9]. Although seemingly simpler than other sequential decision-making problems such as the Markov decision process (MDP) [10], [11] and the partially observable Markov decision process (POMDP) [12], [13], the theory of MABs has long been the subject of research in the field of statistical machine learning and has been applied to solving problems in the fields of recommendation systems, healthcare, and finance [14]. While there are several major formulations of MABs such as stochastic bandits [15] and adversarial bandits [16], depending on how uncertain outcomes are handled, here we focus on stochastic bandits treating both present and future outcomes in the same fashion, i.e. without any discounting. This simplifies the development and is appropriate in settings where it may not be possible to predict in advance how many iterations will be required for an agent to learn an action selection policy.

More formally written, a stochastic MAB problem aims to minimize the cumulative regret, i.e. the sum of difference between the maximum expected outcome and an actual outcome over iterations, via alternately performing two steps: (1) action selection: selecting one option (a.k.a “arm”) from alternatives with unknown outcome distributions; and (2) observation update: updating the statistics for the selected option based on the outcome obtained. In stochastic bandits, it is common to apply the Bayes’ rule to update the posterior of the estimated outcome distribution of the selected option in step (2). For step (1), various action selection methods such as ϵ -greedy, the upper confidence bound (UCB) [15], and the Thompson sampling (TS) [17] have been devised.

In standard MAB problems, only the information about each option itself (e.g. the cumulative value of outcomes and/or the number of times it is selected) is taken into account during action selection. However, it may be possible to minimize the cumulative regret more effectively and find the best option with fewer iterations, by utilizing contextual information associated with each option (e.g. user age in case of web advertisement). Variants of stochastic MABs that incorporate such secondary information are called contextual multi-armed bandits (CMABs) and have been used for dialogue systems, anomaly detection, etc. [14]. In robotic decision-making problems, it is typical to collect and use multiple sources of side information, so we will focus on CMAB problems hereafter. However, conventional action selection strategies for CMABs use heuristics (e.g. the exploration bonus term in UCB) to select options, and these methods usually require many iterations for informed decision-making [3], [9]. Thus, it is desirable to develop an alternative action selection strategy for robotics applications where the number of iterations cannot be too large.

B. Free-Energy Principle and Active Inference

The free-energy principle is a theory in neuroscience that is expected to explain whole-brain mechanisms [18]. According to this principle, biological agents perceive, learn, and act by minimizing free-energy, which is the upper bound of surprise (i.e. $-\log p(o)$, where o is an observation/outcome) observed from the external world, to maintain their homeostasis according to neurally encoded probabilistic models. In neuroscience, this principle has been used to explain many physiological phenomena, such as predictive coding [19].

Active inference is a mathematical framework that applies the free-energy principle to the behavioral norms of biological agents, and has recently begun to attract attention in not only neuroscience and but also robotics [20]. Since future observations are unobservable until an action is executed, in active inference, decision-making agents are assumed to act to minimize a measure called expected free energy (EFE). The reason why this framework is theoretically sound for sequential decision-making problems is that the EFE is composed of value and information gain terms. Thus, by minimizing the EFE, it is possible to naturally balance the tradeoff between exploitation and exploration in a principled non ad-hoc manner. Additionally, active inference can easily incorporate prior information about the probability distribution of outcomes that agents desire to observe, which is called a prior preference (a.k.a “evolutionary prior”). By adjusting this distribution, agent behavior can be varied. Adjustment of the prior preference values corresponds to a type of reward tuning required in typical sequential decision-making problems. Yet, specifying the probability of observed outcomes desired by the agent arguably provides a theoretically intuitive and elegant alternative to specifying numerical reward values over intermediate actions and states.¹

¹There exist studies attempting to estimate an appropriate prior preference in terms of reward functions via inverse reinforcement learning [21], however this type of learning is not the focus of our study.

However, active inference has only recently been considered for multi-armed bandit problems, namely in [8]. Their work has shown that an active inference action selection strategy produces higher performance (i.e. smaller cumulative regrets) than the state-of-the-art methods, especially for stationary MABs when the number of total iterations are relatively small and for dynamic switching MABs. Nevertheless, their study does not focus on contextual bandits, which are arguably even more applicable and interesting for autonomous robotics. Thus, in the following, after reviewing CMAB problems (where in particular outcomes are categorical), we explain how the EFE expression is derived in these problems. We then derive novel EFE approximations to address the fact that the resulting terms are analytically intractable due to the observation likelihoods.

III. METHODOLOGIES

A. Problem Statement

In a contextual multi-armed (CMAB) problem introduced in Section II-A, suppose the total number of options (actions a) taken into account by an autonomous robot is K . These options correspond to the bandit arms. Also suppose that the observed outcome o_k of each option/action a is binary, i.e. $o_k \in \{0, 1\} \forall k \in \{1, \dots, K\}$.² Note, we use the shorthand the notation $a_k \leftrightarrow a = k$.

Since the number of outcomes is binary, the probability ψ_k that each option k outputs a preferred outcome, $o_k = 1$, depends on a hidden linear parameter vector $\vec{\theta}_k$ (unique to each option) and a shared context vector \vec{x} , and is represented by the following sigmoid function (without loss of generality, \vec{x} is assumed common for all options and iterations),

$$\psi_k = p(o_k = 1) = \frac{e^{\vec{\theta}_k^T \vec{x}}}{1 + e^{\vec{\theta}_k^T \vec{x}}} = \frac{1}{1 + e^{-\vec{\theta}_k^T \vec{x}}}. \quad (1)$$

In the case of such a Bernoulli contextual bandit, the expected outcome of each option is also ψ_k so that the cumulative regret is expressed as follows,

$$\text{Regret}(T) = T\psi^* - \sum_{k=1}^K N_T(k)\psi_k, \quad (2)$$

where T is the total number of iterations, ψ^* is the maximum expected outcome (unknown *a priori* for the robot), and $N_T(k)$ represents how many times an option k is selected within T iterations. The goal of CMAB problems is to minimize expected cumulative regret. To do so, the robot needs to estimate the values of the hidden parameter vectors $\vec{\theta}_k \forall k, k \in \{1, \dots, K\}$ in the process of finding an optimal option by iteratively performing the two steps of action selection and measurement update.

B. Expected Free Energy in CMABs

As already mentioned in Sec. II-B, according to the free-energy principle, an agent wants to avoid (i.e. minimize) surprise to maintain homeostasis. The surprise in the case of

²The methodology described in this section is extendable to more than two outcome categories with a few straightforward modifications.

the CMAB problem defined in Sec. III-A can be expressed by marginalizing the following joint distribution,

$$-\log p(o) = -\log \int_{\vec{\theta}_1} \dots \int_{\vec{\theta}_K} p(o, \vec{\theta}_1, \dots, \vec{\theta}_K) d\vec{\theta}_1 \dots d\vec{\theta}_K. \quad (3)$$

However, calculating (3) directly via multiple integral is often analytically intractable, so instead its upper bound (i.e. free-energy) is minimized. Yet, outcomes are unobservable until an option is executed, in the case of active inference, the expected free energy (EFE) is considered. Note, in the following, $\vec{\Theta}$ represents $[\vec{\theta}_1, \dots, \vec{\theta}_K]^T$,

$$\begin{aligned} \text{EFE}(a_k) &= \int_{\vec{\Theta}} q(\vec{\Theta}|a_k) \sum_o p(o|\vec{\Theta}, a_k) \log \frac{q(\vec{\Theta}|a_k)}{p(\vec{\Theta}, o|a_k)} d\vec{\Theta} \\ &\approx \int_{\vec{\Theta}} q(\vec{\Theta}|a_k) \sum_o p(o|\vec{\Theta}, a_k) \log \frac{q(o|a_k)}{p_{ev}(o) \cdot p(o|\vec{\Theta}, a_k)} d\vec{\Theta}, \quad (4) \end{aligned}$$

where $q(\vec{\Theta}|a_k)$ is a proposal distribution that approximates the posterior distribution $p(\vec{\Theta}|o, a_k)$, $p(o|\vec{\Theta}, a_k)$ is an observation likelihood, $q(o|a_k)$ is the marginalization of $q(o, \vec{\Theta}|a_k)$, and $p_{ev}(o)$ is a prior preference. For brevity, by further assuming that the hidden linear parameter vectors are independent each other, (4) can be rewritten as

$$\begin{aligned} \text{EFE}(a_k) &\approx \int_{\vec{\theta}_k} q(\vec{\theta}_k|a_k) \sum_o p(o|\vec{\theta}_k, a_k) \left\{ \log \frac{q(o|a_k)}{p_{ev}(o)} - \log p(o|\vec{\theta}_k, a_k) \right\} d\vec{\theta}_k \\ &= -\mathbb{E}_{q(o|a_k)} \left[\log p_{ev}(o) \right] - \mathbb{E}_{q(o|a_k)} \left[D_{KL}(q(\vec{\theta}_k|o, a_k) || q(\vec{\theta}_k|a_k)) \right]. \quad (5) \end{aligned}$$

The full derivation of (4) and (5) are found in [22]. As can be seen in (5), minimizing EFE value naturally achieves balancing exploitation (the first term) and exploration (the second term). Thus, in CMAB problems with active inference, the robot needs to choose the option minimizing the EFE term in each action selection step,

$$a_{selected} = \arg \min_{a \in |A|} \text{EFE}(a). \quad (6)$$

Yet, as the observation likelihood $p(o|\vec{\theta}_k, a_k)$ is the sigmoid function, the EFE term cannot be computed analytically. Hence, in the following subsections, we introduce the two statistical approximation methods: (1) variational Bayesian importance sampling (VBIS) [23], which has been applied to other probabilistic decision-making under uncertainty problems in robotics [24], [25]; and (2) Laplace approximation [26], which is commonly used in statistical machine learning. Note that, instead of the sigmoid function, a conditional probability table (CPT) [26] could be used to express the observation likelihood and achieve tractable posterior expressions for the EFE (and CMAB inference more generally using other approaches like ϵ -greedy). However, this is generally far less practical, since the number of CPT parameters increases exponentially with the dimension of \vec{x} , whereas this only increases linearly for the sigmoid likelihood.

C. VBIS EFE Approximation

The variational Bayesian importance sampling (VBIS) is a hybrid data fusion method that approximates an analytically intractable posterior distribution in human-robot collaborative sensing when a flexible discrete-continuous function, e.g.

softmax function, is used as an observation likelihood [23]. The main idea behind the VBIS algorithm is that a softmax function can be lower-bounded by a variational Gaussian via the inequality proved in [27] so that if a prior distribution is a Gaussian, a posterior distribution is also approximated as another Gaussian. Here, since the logistic sigmoid function in the EFE expression is a special case of a softmax function with two classes, the following log-sum-exponential (LSE) function is upper bounded as follows

$$\log \left(\sum_{h=0}^1 e^{y_h} \right) \leq \alpha + \sum_{h=0}^1 \frac{y_h - \alpha - \xi_h}{2} + \lambda(\xi_h) \left[(y_h - \alpha)^2 - \xi_h^2 \right] + \log(1 + e^{\xi_h}), \quad (7)$$

where $y_0 = 0$, $y_1 = \vec{\theta}^T \vec{x}$, $\lambda(\xi_h) = \frac{1}{2\xi_h} \left[\frac{1}{1+e^{-\xi_h}} - \frac{1}{2} \right]$. α and ξ_h are free variational parameters which can be adjusted to achieve the tighter bound of the LSE and be iteratively updated within the VBIS procedure. By replacing the denominator of (1) with this bound, the sigmoid function is lower bounded by the variational Gaussian.

$$p(o = j|\vec{\theta}_k) \geq \exp \left(\mathcal{G} + \mathcal{H}^T \vec{\theta}_k - \frac{1}{2} \vec{\theta}_k^T \mathcal{K} \vec{\theta}_k \right), \quad (8)$$

where \mathcal{G} , \mathcal{H} , and \mathcal{K} are

$$\mathcal{G} = \sum_{h=0}^1 \left[\frac{\xi_h}{2} + \lambda(\xi_h)(\xi_h^2 - \alpha^2) - \log(1 + e^{\xi_h}) \right], \quad (9)$$

$$\mathcal{H} = \begin{cases} -\frac{1}{2} \vec{x} + 2\lambda(\xi_1) \cdot \alpha \cdot \vec{x} & \text{if } o = 0, \\ \frac{1}{2} \vec{x} + 2\lambda(\xi_1) \cdot \alpha \cdot \vec{x} & \text{if } o = 1, \end{cases} \quad (10)$$

$$\mathcal{K} = 2\lambda(\xi_1) \cdot \vec{x} \cdot \vec{x}^T. \quad (12)$$

Thus, in our problem, if we assume that the prior distribution of a hidden linear parameter vector $q(\vec{\theta}_k|a_k)$ is approximated as a (multivariate) Gaussian, the joint distribution $q(\vec{\theta}_k|a_k) \cdot p(o|\vec{\theta}_k, a_k)$ in (5) becomes also another unnormalized Gaussian. By normalizing this distribution, the approximated normalization constant $\hat{C}_{vb} = \int_{\vec{\theta}_k} q(\vec{\theta}_k|a_k) \cdot \exp \left(\mathcal{G} + \mathcal{H}^T \vec{\theta}_k - \frac{1}{2} \vec{\theta}_k^T \mathcal{K} \vec{\theta}_k \right) d\vec{\theta}_k$ required for calculating (5) can be obtained from the well-known Gaussian integral. During this process, the VB posterior, which approximates the posterior $p(\vec{\theta}_k|o, a_k)$, is also derived as a side product.

$$\begin{aligned} p(\vec{\theta}_k|o, a_k) &\approx p_{vb}(\vec{\theta}_k|o, a_k) \\ &= \frac{q(\vec{\theta}_k|a_k) \cdot \exp(\mathcal{G} + \mathcal{H}^T \vec{\theta}_k - \frac{1}{2} \vec{\theta}_k^T \mathcal{K} \vec{\theta}_k)}{\int q(\vec{\theta}_k|a_k) \cdot \exp(\mathcal{G} + \mathcal{H}^T \vec{\theta}_k - \frac{1}{2} \vec{\theta}_k^T \mathcal{K} \vec{\theta}_k) d\vec{\theta}_k} \end{aligned} \quad (13)$$

However, the value of \hat{C}_{vb} is smaller than the true normalization constant $C = \int_{\vec{\theta}_k} q(\vec{\theta}_k|a_k) p(o|\vec{\theta}_k, a_k) d\vec{\theta}_k$, the importance sampling (IS) is applied to improve the estimate by following [23].

IS can approximate the expected value of an arbitrary function $f(\vec{\theta}_k)$ over a pdf such as a posterior that is difficult to sample directly, by utilizing a proposal distribution that is roughly similar in shape to the posterior and easy to sample. Here, in particular, the IS approximation is constructed as

$$\langle f(\vec{\theta}_k) \rangle = \langle \vec{\theta}_k \rangle \approx \sum_{s=1}^{N_s} w_s \vec{\theta}_{k,s}, w_s \propto \frac{q(\vec{\theta}_{k,s}|a_k) p(o|\vec{\theta}_{k,s}, a_k)}{p_{prop}(\vec{\theta}_{k,s})}, \quad (14)$$

Algorithm 1 Active inference action selection for CMABs

Input: Context vector \vec{x} , evolutionary prior $p_{ev}(o)$, estimated means and covariances, number of samples for importance sampling N_s

Output: a selected index for action selection

```

1: for each option do
2:   for each outcome do
3:     obtain  $\vec{\mu}_{posterior}$ ,  $\Sigma_{posterior}$ , and  $\hat{C}_{posterior}$  from the VBIS
       or the Laplace approximation algorithm [23], [26]
4:     re-derive  $\mathcal{G}'$ ,  $\mathcal{H}'$ ,  $\mathcal{K}'$  terms from (18)
5:     calculate EFE( $a_k, o$ ) via (16) and (19)
6:   end for
7:   calculate EFE( $a_k$ ) =  $\sum_o$  EFE( $a_k, o$ )
8: end for
9: return a selected index =  $\arg \min_{a \in |A|} \text{EFE}(a)$ 

```

where N_s is the number of samples, w_s is the importance weight for sample s , and $p_{prop}(\vec{\theta}_k)$ is the proposal pdf

$$p_{prop}(\vec{\theta}_k) = \mathcal{N}(\vec{\mu}_{vb}, \Sigma_{prior}) \quad (15)$$

where $\vec{\mu}_{vb}$ is the VB mean and Σ_{prior} is the prior covariance matrix of $q(\vec{\theta}_k|a_k)$. During this procedure, the improved normalization constant estimate \hat{C}_{vbis} can be calculated as $\sum_s w_s / N_s$. Thus, the first term inside the summation over outcome o in the second equation of (5) is rewritten as

$$\log \frac{q(o|a_k)}{p_{ev}(o)} \cdot \int_{\vec{\theta}_k} q(\vec{\theta}_k|a_k) p(o|\vec{\theta}_k, a_k) d\vec{\theta}_k = \left\{ \log \frac{C_{vbis}}{p_{ev}(o)} \right\} \cdot C_{vbis}. \quad (16)$$

Similarly, the second term is approximated as

$$\begin{aligned} &\int_{\vec{\theta}_k} q(\vec{\theta}_k|a_k) p(o|\vec{\theta}_k, a_k) \log p(o|\vec{\theta}_k, a_k) d\vec{\theta}_k \\ &= C_{vbis} \int_{\vec{\theta}_k} p_{vbis}(\vec{\theta}_k|o, a_k) \left(\mathcal{G} + \mathcal{H}^T \vec{\theta}_k - \frac{1}{2} \vec{\theta}_k^T \mathcal{K} \vec{\theta}_k \right) d\vec{\theta}_k. \end{aligned} \quad (17)$$

However, the quadratic term in (17) (i.e. $\mathcal{G} + \mathcal{H}^T \vec{\theta}_k - \frac{1}{2} \vec{\theta}_k^T \mathcal{K} \vec{\theta}_k$) has not yet been processed via the IS step. Therefore, it is required to re-update the lower-bounded variational Gaussian from the following identical equation to avoid underestimating an EFE value.

$$C_{vbis} \cdot \mathcal{N}(\vec{\mu}_{vbis}, \Sigma_{vbis}) = \mathcal{N}(\vec{\mu}_{prior}, \Sigma_{prior}) \cdot \mathcal{N}(\vec{\mu}_{sigm}, \Sigma_{sigm}) \quad (18)$$

After re-deriving \mathcal{G}' , \mathcal{H}' , and \mathcal{K}' terms of $\mathcal{N}(\vec{\mu}_{sigm}, \Sigma_{sigm})$, (17) is further transformed as follows.

$$\begin{aligned} (17) &= C_{vbis} \cdot \mathbb{E}_{p_{vbis}(\vec{\theta}_k|o, a_k)} \left[\mathcal{G}' + \mathcal{H}'^T \vec{\theta}_k - \frac{1}{2} \vec{\theta}_k^T \mathcal{K}' \vec{\theta}_k \right] \\ &= C_{vbis} \left(\mathcal{G}' + \mathcal{H}'^T \vec{\mu}_{vbis} - \frac{1}{2} \left(\text{Tr}(\mathcal{K}' \Sigma_{vbis}) + \vec{\mu}_{vbis}^T \mathcal{K}' \vec{\mu}_{vbis} \right) \right) \end{aligned} \quad (19)$$

After subtracting (19) from (16), and sum the EFE(a_k, o) term over possible outcomes o , the EFE(a_k) is computed. Then, by following (6), an active inference robot finalizes the option in an action selection step. Algorithm 1 summarizes the process of active inference action selection for CMABs with the VBIS EFE approximation.

D. Laplace EFE Approximation

In statistical machine learning, the Laplace approximation is commonly used to approximate a continuous probability density function as a Gaussian distribution [26]. This uses the second-order approximation of the vector Taylor expansion

of a logarithmic function whose gradient is $\vec{0}$. Particularly, in the approximation of the EFE value in (5), a function $g(\vec{\theta}_k)$ is defined as $q(\vec{\theta}_k|a_k)p(o|\vec{\theta}_k, a_k)$, and the following logarithmic function is used.

$$\log g(\vec{\theta}_k) \approx \log g(\vec{\theta}_k^{(0)}) + \sum_{r=1}^n (\theta_{k,r} - \theta_{k,r}^{(0)}) \frac{\partial g(\vec{\theta}_k^{(0)})}{\partial \theta_{k,r}} + \frac{1}{2} \left\{ \sum_{r=1}^n (\theta_{k,r} - \theta_{k,r}^{(0)}) \frac{\partial \log g(\vec{\theta}_k^{(0)})}{\partial \theta_{k,r}} \right\}^2, \quad (20)$$

where $\vec{\theta}_k^{(0)}$ satisfies $\nabla \log g(\vec{\theta}_k^{(0)}) = \vec{0}$ and n is the length of $\vec{\theta}_k$. However, since it is analytically intractable to find $\vec{\theta}_k^{(0)}$, $\vec{\theta}_{k,MAP}$ is computed via Newton's method [28]. Also, the third term in (20) is rewritten as

$$\left\{ \sum_{r=1}^n (\theta_{k,r} - \theta_{k,r}^{(0)}) \frac{\partial \log g(\vec{\theta}_k^{(0)})}{\partial \theta_{k,r}} \right\}^2 = \left\{ (\vec{\theta}_k - \vec{\theta}_k^{(0)})^T \nabla \log g(\vec{\theta}_k^{(0)}) \right\}^2, \quad (21)$$

therefore, (20) reduces to

$$\log g(\vec{\theta}_k) \approx \log g(\vec{\theta}_{k,MAP}) + \frac{1}{2} (\vec{\theta}_k - \vec{\theta}_{k,MAP})^T H \left[\log g(\vec{\theta}_{k,MAP}) \right] (\vec{\theta}_k - \vec{\theta}_{k,MAP}), \quad (22)$$

where H is the Hessian. Thus, if we define $A = -H$ and by taking the logarithm from (22),

$$g(\vec{\theta}_k) \approx g(\vec{\theta}_{k,MAP}) \cdot \exp \left(-\frac{1}{2} (\vec{\theta}_k - \vec{\theta}_{k,MAP})^T A (\vec{\theta}_k - \vec{\theta}_{k,MAP}) \right), \quad (23)$$

and the normalization constant $C_{laplace}$ is computed as

$$C_{laplace} = g(\vec{\theta}_{k,MAP}) \cdot \frac{(2\pi)^{\frac{n}{2}}}{|A|^{\frac{1}{2}}}. \quad (24)$$

After following the same procedures as with the VBIS described in Algorithm 1 (i.e. calculating the Laplace posterior $\mathcal{N}(\vec{\mu}_{laplace}, \Sigma_{laplace})$ and re-deriving \mathcal{G}' , \mathcal{H}' , and \mathcal{K}' terms of $\mathcal{N}(\vec{\mu}_{sigm}, \Sigma_{sigm})$), the EFE(a_k) of each option k is calculated, and then the best option can be selected.

IV. SIMULATION STUDY

A. Motivating Problem Scenario

NASA recently launched the Concepts for Ocean worlds Life Detection Technology (COLDTech) program to research and develop software for 'fail-operational' autonomous robotic landers that will be dispatched to icy solar system moons like Europa and Enceladus [2]. Currently, the concept of operations calls for a robotic lander to perform basic surface tasks such as digging, imaging, moving objects, retrieving/analyzing samples with various instruments, and downlinked selected science data. As shown in Fig. 1, a key capability will be autonomous selection of surface areas to prioritize for science investigation. Although human experts can specify preferences for possible areas in advance, the unknown dynamic nature of icy moon environments can cause science priorities to shift unexpectedly, e.g. due to plume eruptions or changes in surface conditions. Such priorities must be considered in light of contextual information pertaining to overall mission progress, lander state, and engineering factors.

In the following, we consider proof of concept CMAB simulations for a robotic lander which has been given the task of collecting an ice sample specimen from one of K non-overlapping designated reachable areas within the landing site. The candidate sampling sites are selected by scientists based on imaging data collected immediately after landing. Since a sampling action consumes a considerable amount of resources (e.g. energy and time for grinding, drilling, scooping, etc.), the lander cannot just deploy a resource-intensive manipulator in a haphazard way. It is necessary to infer in advance which areas can yield scientifically valuable information, i.e. $o_k = 1$, using a laser spectrometer similar to those used on Martian rovers [29], which can be more easily and cheaply deployed than the manipulator tools. Panoramic images of the landing area taken onboard are also processed by a learning-based computer vision algorithm to produce a C -dimensional binary context feature vector \vec{x} , to encode various factors such as sun direction, surface reflectance, etc. [30], [31]. The autonomous site selection problem can thus be formulated as a CMAB, where the goal is to find the best site k to sample with the manipulator, using observations $o_k = 1/0$ (interesting/uninteresting science returns) from selected spectrometer target sites a_k and context \vec{x} .

B. Simulation Setup and Results

As described in Sec. III-A, CMAB problems require alternately executing two steps: action selection (which site to illuminate with laser spectrometer) and observation update (process spectrometer returns). In this simulation study, the following solution approaches are considered and compared in an extensive Monte Carlo study: (i) best possible action (site) selection, using an offline oracle (needed to compute regret); (ii) ϵ -greedy (where $\epsilon = 0.25$ was found to work best after initial trials); (iii) upper confidence bound (UCB); (iv) logistic Thompson sampling (LTS); and (v) active inference (AI). The action selection methods for (iv) and (v) are paired with VBIS and Laplace approximations for the observation updates. In all simulations, 100 Monte Carlo (MC) runs are performed, and the number of iterations T in each MC run is set to 10^2 , which is small compared to common MAB algorithm benchmarks [8] and reflects a practical upper limit for robotic lander sensor deployment. Thus, there is no need to be too concerned about the suboptimal asymptotic behavior of active inference agents on very long time scales (e.g. 10^5 iterations) as experimentally showed in [8]. The true hidden linear parameters $\vec{\theta}_k$ for each candidate search site k were randomly generated from a multivariate Gaussian distributions, where the mean is sampled from a uniform distribution of 0 to 1, and the covariance matrix is generated using the result of a singular value decomposition on the Gram matrix of a matrix with uniformly randomly distributed entries. The search context vector \vec{x} was randomly generated assuming that each element takes a binary value with uniform probability, and the same value is used for all sites. To compare how each approach performs with respect to the difficulty of CMAB, the number of areas K and search context elements C , and the evolutionary prior $p_{ev}(o)$, were

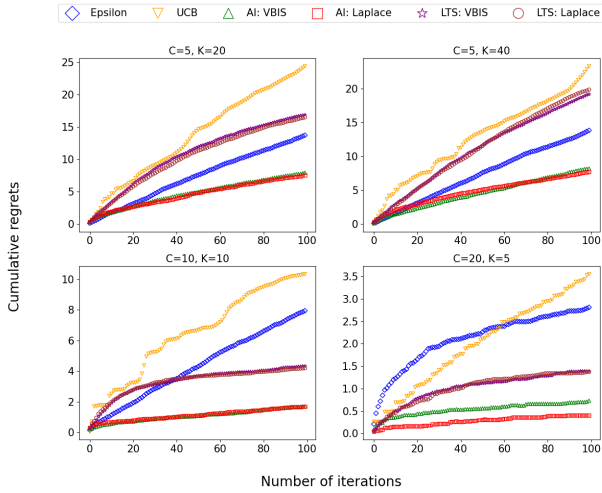


Fig. 2. Comparison of cumulative regrets among different approaches. For each subplot, the cumulative regret is the average of 100 simulations. Prior preference of active inference agents is set as $p_{ev}(o) = [0.001, 0.999]^T$.

varied, such that $K \in \{5, 10, 20, 40\}$, $C \in \{5, 10, 20\}$, and $p_{ev}(o) \in \{[0.001, 0.999]^T, [0.4, 0.6]^T\}$, i.e. in total, $24 \times 7 = 168$ sets of 100 MC sims were run. This setup produces difficult CMABs since there can be multiple “good” (large ψ) sites, which makes it hard to distinguish the best site with highest chance of good science returns.

Results: The cumulative regret (2) was used to compare the performance of all approaches. Fig. 2 shows the characteristic results, including those obtained when the easiest (bottom right; the most amount of information and the fewest number of options) and most difficult (top right; lowest amount of information and most options) conditions are used. As can be seen from this figure, as the difficulty of the CMAB problem increases (counterclockwise from the bottom right to top right), the values of cumulative regrets tend to increase regardless of which action selection strategy is used, nevertheless the both AI VBIS and AI Laplace robots greatly outperforms the others (Note: the slight performance differences between the AI VBIS and AI Laplace robots are mainly from the accuracy of the posterior distribution approximation). This is owing to the fact that, as shown in (5), AI can rigorously evaluate the uncertainty of each option and quickly find large ψ options, and prefers those options under the influence of the prior preference $p_{ev}(o)$. On the other hand, when the prior preference is non-reward seeking and favors more uncertain outcome distributions (i.e. $p_{ev}(o) = [0.4, 0.6]^T$), the AI robots perform as comparable as other action selection strategies as presented in Fig. 3. In such a case, in order to assess the performance of the AI robots, other evaluation metric such as the Kullback-Leibler divergence may be utilized (i.e. the extent to which the outcome distribution obtained by executing the option preferred by the AI robots are aligned with the prior preference). In terms of computation cost, although the cost of AI VBIS approach is higher than other strategies due to importance sampling, this can be significantly reduced by computing sample weights in parallel. With regard to AI Laplace, while

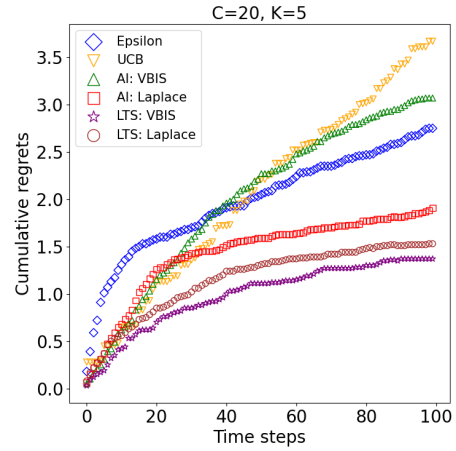


Fig. 3. Cumulative regrets among different approaches in case of $C=20$ and $K=5$. Prior preference is non-reward seeking and favors more uncertain outcome distributions, i.e. $p_{ev}(o) = [0.4, 0.6]^T$.

the average single action selection time (0.010 sec in case of $K = 10$, $C = 10$) is again higher than that of ϵ -greedy ($3.42E-06$ sec in the same condition), still the computation cost may not be problematic since the number of total iterations is small enough (at most 10^2) and the performance of the AI Laplace robot outperforms the baseline methods. Even in the most difficult case ($K = 40$, $C = 20$), AI Laplace takes only 0.26 sec on average, compared to $5.48E-06$ sec for ϵ -greedy.

V. CONCLUSIONS

In this paper, we applied active inference as an action selection strategy for contextual multi-armed bandit (CMAB) problems. We showed how to derive the expected free energy (EFE) in categorical outcome CMAB problems and introduced variational and Laplace approximation algorithms to cope with analytically intractable EFE terms. The proposed active inference strategy was demonstrated in a simulation for an autonomous remote science site selection scenario, and its cumulative regret was compared with the ones of other strategies. It was found that when the prior preference is reward-seeking, active inference brings significant advantages, due to its ability to evaluate option uncertainty and bias towards preferred outcomes.

The primary objective of this study was to validate the feasibility and performance of active inference strategy in practical CMAB problems using proof of concept simulations that reflect key aspects of the motivating remote space science exploration problem. Future work will confirm the effectiveness of this strategy using OceanWATERS, [32], a testbed for the Europa lander recently released by NASA. In this work, the value of prior preferences was determined and fixed *a priori*, although adjusting this value according to search context, the regret analysis of active inference action selection strategies and results obtained from the actual action execution are also future tasks.

REFERENCES

- [1] C. B. Phillips and R. T. Pappalardo, "Europa Clipper Mission Concept: Exploring Jupiter's Ocean Moon," *EOS Transactions*, vol. 95, no. 20, pp. 165–167, May 2014.
- [2] J. McMahon, N. Ahmed, M. Lahijanian, P. Amorese, T. Deka, K. Muvvala, T. Slack, and S. Wakayama, "Expert-informed autonomous science planning for in-situ observations and discoveries," in *2022 IEEE Aerospace Conference (AERO)*, 2022, pp. 1–11.
- [3] L. Li, W. Chu, J. Langford, and R. E. Schapire, "A contextual-bandit approach to personalized news article recommendation," in *WWW*. ACM, 2010, pp. 661–670.
- [4] S. McGuire, P. M. Furlong, T. Fong, C. Heckman, D. Szafir, S. J. Julier, and N. Ahmed, "Everybody needs somebody sometimes: Validation of adaptive recovery in robotic space operations," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1216–1223, 2019.
- [5] L. Chan, D. Hadfield-Menell, S. Srinivasa, and A. Dragan, "The assistive multi-armed bandit," in *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '19. IEEE Press, 2019, p. 354–363.
- [6] K. Friston, F. Rigoli, D. Ognibene, C. Mathys, T. FitzGerald, and G. Pezzulo, "Active inference and epistemic value," *Cognitive neuroscience*, 02 2015.
- [7] R. Kaplan and K. J. Friston, "Planning and navigation as active inference," *Biol. Cybern.*, vol. 112, no. 4, p. 323–343, aug 2018. [Online]. Available: <https://doi.org/10.1007/s00422-018-0753-2>
- [8] D. Markovic, H. Stojic, S. Schwobel, and S. Kiebel, J., "An empirical evaluation of active inference in multi-armed bandits," *Neural Networks; 2021 Special Issue on AI and Brain Science: AI-powered Brain Science*, vol. 144, p. 229–246, may 2021.
- [9] V. Kuleshov and D. Precup, "Algorithms for multi-armed bandit problems," *Journal of Machine Learning Research*, vol. 1, 02 2014.
- [10] R. BELLMAN, "A markovian decision process," *Journal of Mathematics and Mechanics*, vol. 6, no. 5, pp. 679–684, 1957. [Online]. Available: <http://www.jstor.org/stable/24900506>
- [11] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.
- [12] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artif. Intell.*, vol. 101, no. 1–2, p. 99–134, may 1998.
- [13] H. Kurniawati, "Partially observable markov decision processes and robotics," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, no. 1, pp. 253–277, 2022.
- [14] D. Bouneffouf, I. Rish, and C. Aggarwal, "Survey on applications of multi-armed and contextual bandits," in *2020 IEEE Congress on Evolutionary Computation (CEC)*, 2020, pp. 1–8.
- [15] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Mach. Learn.*, vol. 47, no. 2–3, p. 235–256, may 2002. [Online]. Available: <https://doi.org/10.1023/A:1013689704352>
- [16] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. Schapire, "Gambling in a rigged casino: The adversarial multi-armed bandit problem," in *Proceedings of IEEE 36th Annual Foundations of Computer Science*, 1995, pp. 322–331.
- [17] W. R. Thompson, "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples," *Biometrika*, vol. 25, pp. 285–294, 1933.
- [18] K. Friston, "The free-energy principle: A unified brain theory?" *Nature Reviews Neuroscience*, vol. 11, no. 2, pp. 127–138, 2010.
- [19] K. Friston and S. Kiebel, "Predictive coding under the free-energy principle," *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, vol. 364, pp. 1211–21, 06 2009.
- [20] P. Lanillos, C. Meo, C. Pezzato, A. A. Meera, M. Baioumy, W. Ohata, A. Tschantz, B. Millidge, M. Wisse, C. L. Buckley, and J. Tani, "Active inference in robotics and artificial agents: Survey and challenges," *CoRR*, vol. abs/2112.01871, 2021. [Online]. Available: <https://arxiv.org/abs/2112.01871>
- [21] J. Y. Shin, C. Kim, and H. J. Hwang, "Prior preference learning from experts: Designing a reward with active inference," *Neurocomputing*, vol. 492, pp. 508–515, 2022.
- [22] S. Wakayama and N. Ahmed, "Active inference for autonomous decision-making with contextual multi-armed bandits," *CoRR*, vol. abs/2209.09185, 2022. [Online]. Available: <https://arxiv.org/abs/2209.09185>
- [23] N. R. Ahmed, E. M. Sample, and M. Campbell, "Bayesian multicategorical soft data fusion for human–robot collaboration," *IEEE Transactions on Robotics*, vol. 29, no. 1, pp. 189–206, 2013.
- [24] L. Burks, I. Loeffgren, and N. Ahmed, "Optimal continuous state pomdp planning with semantic observations: A variational approach," *IEEE Transactions on Robotics*, vol. 35, pp. 1488–1507, 2019.
- [25] S. Wakayama and N. Ahmed, "Probabilistic semantic data association for collaborative human-robot sensing," *CoRR*, vol. abs/2110.09621, 2021. [Online]. Available: <https://arxiv.org/abs/2110.09621>
- [26] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.
- [27] G. Bouchard, "Efficient bounds for the softmax function and applications to approximate inference in hybrid models," in *Proc. Neural Inf. Process. Syst. Workshop Approx. Bayesian Inference Continuous/Hybrid Syst.*, 2007.
- [28] A. Galantai, "The theory of newton's method," *Journal of Computational and Applied Mathematics*, vol. 124, pp. 25–44, 2000.
- [29] J. Manne, T. Q. Bui, and C. R. Webster, "Determination of foreign broadening coefficients for methane lines targeted by the tunable laser spectrometer (tls) on the mars curiosity rover," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 191, pp. 59–66, 2017.
- [30] L. Matthies, M. Maimone, A. Johnson, Y. Cheng, R. Willson, C. Villalpando, S. Goldberg, A. Huertas, A. Stein, and A. Angelova, "Computer vision on mars," *International Journal of Computer Vision*, vol. 75, no. 1, pp. 67 – 92, October 2007.
- [31] L. Palafox, C. Hamilton, S. Scheidt, and A. Alvarez, "Automated detection of geological landforms on mars using convolutional neural networks," *Computers & Geosciences*, vol. 101, 01 2017.
- [32] Ocean worlds autonomy testbed for exploration research & simulation (oceanwaters). [Online]. Available: https://github.com/nasa/ow_simulator