## A Bayesian Approach to Learning Bandit Structure in Markov Decision Processes

Kelly W. Zhang<sup>1</sup>

**Omer Gottesman**<sup>2</sup>

Finale Doshi-Velez<sup>1</sup>

<sup>1</sup>Department of Computer Science, Harvard University <sup>2</sup>Department of Computer Science, Brown University

#### Abstract

In the reinforcement learning literature, there are many algorithms developed for either Contextual Bandit (CB) or Markov Decision Processes (MDP) environments. However, when deploying reinforcement learning algorithms in the real world, even with domain expertise, it is often difficult to know whether it is appropriate to treat a sequential decision making problem as a CB or an MDP. In other words, do actions affect future states, or only the immediate rewards? Making the wrong assumption regarding the nature of the environment can lead to inefficient learning, or even prevent the algorithm from ever learning an optimal policy, even with infinite data. In this work we develop an online algorithm that uses a Bayesian hypothesis testing approach to learn the nature of the environment. Our algorithm allows practitioners to incorporate prior knowledge about whether the environment is that of a CB or an MDP, and effectively interpolate between classical CB and MDP-based algorithms to mitigate against the effects of misspecifying the environment. We perform simulations and demonstrate that in CB settings our algorithm achieves lower regret than MDP-based algorithms, while in non-bandit MDP settings our algorithm is able to learn the optimal policy, often achieving comparable regret to MDP-based algorithms.

## **1 INTRODUCTION**

Sequential decision making problems are commonly analyzed using two different frameworks: Contextual Bandits (CB) and Markov Decision Processes (MDP). In contextual bandit environments, it is assumed that states evolve independently of actions selected by the algorithm, whereas for MDP environments, action selections affect state transition probabilities.<sup>1</sup> As a consequence in MDP environments, action selections may have long term effects on which states are visited in the future; therefore an algorithm designed for a bandit environment, which does not take such long term effects into account, may never be able to learn an optimal policy. On the other hand, if an environment is known satisfy CB assumptions, this restriction of the problem can be leveraged for more efficient learning, and therefore algorithms designed for bandit environments learn more effectively, as long as the CB assumptions hold. To illustrate this, in Figure 1 we compare the performance of CB and MDP posterior sampling algorithms in both a contextual bandit and an MDP setting, and demonstrate that each algorithm significantly outperforms its competitor in the setting it was designed for.

Reinforcement learning (RL) algorithms are typically developed assuming the environment falls strictly within one of these two frameworks. However, when deploying RL algorithms in the real world, it is often unknown which one of these frameworks one should assume. For example, consider a mobile health application aimed to help users increase their step count [Liao et al., 2020]. A few times a day, the app's algorithm uses the user's state information (e.g. recent app engagement, local weather, time of day, etc.) to decide whether or not to send the user a message encouraging them to take a walk. It is not immediately clear whether to use a contextual bandit or an MDP algorithm in this problem setting. It may be that sending notification strongly affects the future state of the user-for example, notifications that annoy the user could lead them to disengage; if this is a possibility, an MDP algorithm is favored, because a CB algorithm will not be able to learn the optimal policy. On the other hand, it may be that users' states are not greatly affected by notification-for example, users' responsiveness may depend purely on factors such as weather or their work schedule, and past messages do not play a role. In this setting, a CB algorithm could perform better, and making a general MDP assumption can lead to inefficient

<sup>&</sup>lt;sup>1</sup>Note that CB environments is simply a special case of MDPs.



Figure 1: Above we plot the cumulative regret of a contextual bandit posterior sampling algorithm and a MDP posterior sampling algorithm in two environments (finite horizon 100, 6 states). The MDP environment is the river swim environment (see Figure 2). The contextual bandit environment that is identical to the MDP environment, except the state transition probabilities are uniform over all states. The error bars are standard errors for the estimates (100 repetitions).

learning and incurring greater regret.

In this work, we develop the *Bayesian Hypothesis Test*ing Reinforcement Learning (BHT-RL) algorithm, which is an online algorithm for the finite-horizon episodic setting, which utilizes Bayesian Hypothesis Testing to learn whether the environment is that of a CB or a classical MDP. Practitioners specify the prior probability on the environment being a bandit, as well as choose two RL algorithms: one CB algorithm and one MDP algorithm. At the start of each episode, BHT-RL selects which algorithm to use, CB vs. MDP, according to the posterior probability of the environment being a bandit. The choice of prior probability allows the practitioner to choose the amount to "regularize" the MDP algorithm through the amount of evidence needed to favor using the MDP algorithm over the CB algorithm. We empirically demonstrate that in terms of regret minimization, BHT-RL outperforms CB and MDP algorithms in their respective misspecified environments, and that BHT-RL is often competitive with both CB and MDP algorithms each in the environments they were designed for. Moreover, BHT-RL results in significantly better regret minimization empirically in both contextual bandit and MDP environments compared prior work on learning bandit structure in MDPs [Zanette and Brunskill, 2018, 2019]. Additionally, when posterior sampling algorithms are chosen for the contextual bandit and MDP algorithms, then the BHT-RL algorithm can be interpreted as posteroir sampling (PS) with an additional prior that can up-weight the prior probability that environment is a bandit; thus, the Bayesian regret bounds for standard MDP-PS apply [Osband et al., 2013].

## 2 RELATED WORK

An open problem in RL theory is understanding how sample complexity depends on the planning horizon (in infinite horizon problems, this is the discount factor) [Jiang and Agarwal, 2018]. Jiang et al. [2015] showed that when learning the optimal policy from data in MDP problems, longer planning horizons increase the size of the set of policies one searches over. As a result, often when one has small amounts of data it is better to use a smaller planning horizon than the evaluation horizon as a method of regularization to prevent overfitting to the data [Arumugam et al., 2018]. However, using a shorter planning horizon can also prevent algorithms from ever learning the optimal policy-for example, a CB algorithm, which has a planning horizon of 1, will never be able to learn the optimal policy in most MDP environments. Since the BHT-RL algorithm interpolates between CB and MDP algorithms, it can be viewed as a regularized version of standard MDP algorithms which reduces the planning horizon, i.e., using a CB algorithm, when the evidence that the environment is an MDP is low.

There have been several works developing low regret algorithms for both CB and MDP environments. Jiang et al. [2017] develop the OLIVE algorithm and prove bounds for it in a variety of sequential decision making problems when the the Bellman rank is known. However, since the Bellman rank of problems is generally unknown in real world problems, OLIVE cannot be used in practice. Zanette and Brunskill [2018] develop the UBEV-S RL algorithm, which they prove has near optimal regret in the MDP setting and has regret that scales better than that of OLIVE in the CB setting [Zanette and Brunskill, 2018]. Later, Zanette and Brunskill [2019] developed the EULER algorithm, which improves upon UBEV-S, and has optimal regret bounds in both the contextual bandit and MDP settings [Zanette and Brunskill, 2019]. Both UBEV-S and EULER are upper confidence bound based methods that construct confidence bounds for the next timestep reward and future value, and then execute the most optimistic policy within those bounds. Although there have only been Bayesian regret bounds (in

contrast to frequentist regret bounds) proven for posterior sampling on MDPs [Osband et al., 2013], it has been shown in previous work that posterior sampling RL algorithms generally outperform confidence bound based algorithms empirically [Osband et al., 2013, Osband and Van Roy, 2017, Russo and Van Roy, 2014], and in our experiment we see that BHT-RL outperforms both UBEV-S and EULER.

For BHT-RL, we pool the state transition counts for different actions in the same state together. Asmuth et al. [2009] call this approach the tied Dirichlet model. However, they also assume that the experimenter has apriori knowledge and chooses before the study is run whether to assume the tied or regular Dirichlet model on the transition probabilities. In contrast, we will aim to *learn* whether in each state it is better to use the tied Dirichlet model or the standard one. To do this we will use Bayesian hypothesis testing [Berger, 2012]. Bayesian hypothesis testing is related to Bayesian model selection because the posterior probabilities of the null versus the alternative models are a function of the Bayes factor, which is used in model selection to compare the relative plausibilities of two different models or hypotheses.

## **3** BAYESIAN HYPOTHESIS TESTING RL

#### 3.1 PROBLEM SETTING

We define random variables for the states  $S_t \in S$ , random variables for action selections  $A_t \in A$ , and rewards  $R_t \in \mathbb{R}$ . We also define  $\theta$ , which parameterizes the environment, i.e., given  $\theta$  we know the expected rewards  $E_{\theta}[R_t|A_t =$  $a, S_t = s$  and the transition probabilities  $P_{\theta}(S_{t+1} = s' | S_t = s')$  $s, A_t = a$ ). We assume a finite-horizon episodic setting, so the data collected is made up of episodes each of length H. For example, for the  $k^{th}$  episode, we have the data  $(A_{t_k+h}, S_{t_k+h}, R_{t_k+h})_{h=1}^H, \text{ where } t_k := kH. \text{ We define } \mathcal{H}_{t_k} = \{(S_{t_k+h}, A_{t_k+h}, R_{t_k+h})_{h=1}^H\}_{k'=0}^{k-1} = \{(A_t, S_t, R_t)\}_{t=1}^{t_k} \text{ to be his-}$ tory at time  $t_k$ . Note that we define our policies to be  $\pi_k$ to be  $\sigma(\mathcal{H}_{t_k})$ -measurable functions from  $\mathcal{S} \times [1:H]$  to  $|\mathcal{A}|$ dimensional simplex. So, our actions  $A_{t_k+h} \sim \pi_k(S_{t_k+h},h)$ are chosen according to the policy. Note that the policy takes the time-step in the episode, h, as an input because in the finite horizon setting the optimal policy can change depending on the timestep in the episode.

#### **3.2 ALGORITHM DEFINITION**

For our Bayesian Hypothesis Testing method we define the following null and alternative hypotheses. Throughout, we focus on the discrete state setting, but these hypotheses and the BHT-RL method could be generalized to continuous states, when one has a model for the transition probabilities. **Null hypothesis**  $H_0$ : Action selections do not affect transition probabilities, i.e., for all  $a \in A$ ,  $s, s' \in S$ ,

$$P_{\theta}(S_{t+1} = s' | S_t = s, A_t = a) = P_{\theta}(S_{t+1} = s' | S_t = s)$$

Under the null hypothesis we model our data as generated by the following process:

- For each s ∈ S we draw θ̃s from a prior distribution over the transitions. For example, we will use θ̃s ~ Dirichlet(α) for some S-dimensional vector α with positive entries in our derivations and simulations.
- For all  $t \in [1: T]$  such that  $S_t = s$ , we have that  $S_{t+1} \sim$ Categorical $(\tilde{\theta}_s)$ .

Alternative hypothesis  $H_1$ : Action selections do affect transition probabilities, i.e., for some  $a, a' \in A$ ,  $s, s' \in S$ ,

$$P_{\theta}(S_{t+1} = s' | S_t = s, A_t = a) \neq P_{\theta}(S_{t+1} = s' | S_t = s, A_t = a')$$

Under the alternative hypothesis we model our data as generated by the following process:

- For each s ∈ S and each a ∈ A we draw θ
  <sub>s,a</sub> from a prior distribution over the transition probabilities. For example, θ
  <sub>s,a</sub> ~ Dirichlet(α).
- For all  $t \in [1: T]$  such that  $S_t = s$  and  $A_t = a$ , we have that  $S_{t+1} \sim \text{Categorical}(\theta_{s,a})$ .

BHT-RL, as defined in Algorithm 1, requires one to choose a contextual bandit algorithm (denoted  $\pi^{\text{CB}}$ ), an MDP based algorithm (denoted  $\pi^{\text{MDP}}$ ), a generative model for the transition probabilities under both hypotheses, and a prior probability over the hypotheses  $P(H_0)$  and  $P(H_1) = 1 - P(H_0)$ . At the start of each episode, BHT-RL selects which algorithm to use,  $\pi^{\text{CB}}$  vs.  $\pi^{\text{MDP}}$ , according to the posterior probability of the null hypothesis. Note that if we set the prior probability of the null  $P(H_0)$  to 1 the BHT-RL algorithm is equivalent to  $\pi^{\text{CB}}$  and when setting  $P(H_0)$  to 0 the BHT-RL algorithm is equivalent to  $\pi^{\text{MDP}}$ . Practically, for someone utilizing the algorithm, the choice of  $P(H_0)$  depend on how likely they think that the environment is that of a bandit, based on domain knowledge. Then given we have run *k* episodes already we can compute the posterior probabilities for the hypotheses,  $P(H_0|\mathcal{H}_{t_k})$  and  $P(H_1|\mathcal{H}_{t_k})$ .

$$\begin{split} P(H_0 | \mathcal{H}_{t_k}) &= \frac{P(H_0, \mathcal{H}_T)}{P(\mathcal{H}_{t_k})} \\ \frac{P(\mathcal{H}_{t_k} | H_0) P(H_0)}{P(\mathcal{H}_{t_k} | H_0) P(H_0) + P(\mathcal{H}_{t_k} | H_1) P(H_1)} \end{split}$$

The term  $\frac{P(\mathcal{H}_{t_k}|H_1)}{P(\mathcal{H}_{t_k}|H_0)}$  above is the Bayes factor. See Appendix A for how we compute the Bayes factor for Dirichlet transition priors.

Note that Bayesian Hypothesis testing approach can be used with any choice of contextual bandit algorithm and

Algorithm 1: Bayesian Hypothesis Testing Reinforcement Learning (BHT-RL)

<b>Input:</b> CB algorithm $\pi^{CB}$ ; MDP algorithm $\pi^{MDP}$ ; Prior
probability $P(H_0)$ ; Generative models for state
transitions under $H_0$ and $H_1$ respectively.
for <i>episodes</i> $k = 0, 1, 2,$ do
Sample indicator of generative model
$B_k \sim \text{Bernoulli}(P(H_0 \mathcal{H}_{t_k}))$
if $B_k = 1$ then
Let $\pi_k = \pi_k^{\text{CB}}$
else
Let $\pi_k = \pi_k^{\text{MDP}}$
for timesteps $h = 1, 2,, H$ do
Sample and apply action $A_t \sim \pi_k(S_{t_k+h}, h)$
Observe $R_{t_k+h}$ and $S_{t_k+h+1}$
end
Update both $\pi_k^{\text{CB}}$ and $\pi_k^{\text{MDP}}$ with data
$\{S_{t_k+h}, A_{t_k+h}, R_{t_k+h}\}_{h=1}^H$ observed in the episode.
end

MDP based algorithm; the generative model for transition probabilities is only used to compute posterior probability  $P(H_0|\mathcal{H}_{t_k})$ . If one chooses posterior sampling methods for the CB and MDP algorithms, then BHT-RL can be interpreted as posterior sampling with a hierarchical prior. Under posterior sampling, a prior is put on the parameters of the environment  $\theta$ . The policy for that episode is selected by first sampling  $\tilde{\theta} \sim Q(\cdot|\mathcal{H}_{t_k})$ , where  $Q(\cdot|\mathcal{H}_{t_k})$  is the posterior distribution over  $\theta$ . Then the policy for the episode  $\pi_k$  is chosen to be the optimal policy for environment  $\tilde{\theta}$ . When using BHT-RL with posterior sampling CB and MDP algorithms, we have that  $\pi_k^{\text{CB}}$  is the optimal policy for  $\tilde{\theta} \sim Q(\cdot|\mathcal{H}_{t_k}, H_0)$ , the posterior distribution of Q given that the null hypothesis  $H_0$  is true. Similarly,  $\pi_k^{\text{MDP}}$  is the optimal policy for  $\tilde{\theta} \sim Q(\cdot|\mathcal{H}_{t_k}, H_1)$ .

# 3.3 LEVERAGING ENDOGENOUS AND EXOGENOUS FEATURES

In many sequential decision making problems, we may know with certainty that the transition probabilities for some state variables do not change depending on the choice of action. This separation into endogenous and exogenous features allows us to learn more efficiently by only testing the null hypothesis for the features which may potentially be affected by actions. For example, in a mobile health setting, we may be uncertain as to whether or not sending a message affects future user engagement, but we do not believe that our messages can affect the future weather. Formally, we assume that the state space S can be decomposed into two parts:  $S = \mathcal{X} \times \mathcal{Z}$ , where  $\mathcal{X}$  includes all known exogenous variables and  $\mathcal{Z}$  includes all potentially endogenous variables. We assume that  $S_t = [X_t, Z_t]$  for  $X_t \in \mathcal{X}$  and  $Z_t \in \mathcal{Z}$  and that for all  $S_t, S_{t+1}$  and  $A_t$ ,

$$P_{\theta}(S_{t+1}|S_t, A_t) = P_{\theta}(X_{t+1}, Z_{t+1}|X_t, Z_t, A_t)$$
$$= P_{\theta}(X_{t+1}|X_t)P_{\theta}(Z_{t+1}|Z_t, A_t)$$

For the reward model, we assume that  $E_{\theta}[R_t|S_t,A_t]$ , the expected reward in a given state and action, is affected by both  $X_t$  and  $Z_t$ . Note that our assumptions differ from other recent works which decompose the state space into endogenous and exogenous components, because we still assume that the exogenous states  $X_t$  can affect the value of the expected reward [Dietterich et al., 2018, Chitnis and Lozano-Pérez, 2020].

Our BHT-RL approach takes advantage of this decomposition to endogenous vs. exogenous feature decomposition by only performing Bayesian hypothesis testing regarding the transition probabilities of sub-states  $Z_t$ , which are potentially endogenous. Performing Bayesian hypothesis testing only on the subset of potentially endogenous states leads to more efficient learning of whether the environment is an CB vs. MDP, compared to performing hypothesis testing regarding the transition probabilities for the entire state  $S_t$ . Moreover, this decomposition can allow BHT-RL to easily scale to much large state spaces when a large number of state features are already known to be exogenous.

#### 3.4 REGRET GUARANTEES

We now define regret in the episodic setting. We first define the value function, which is the expected value of following some policy  $\pi$  during an episode:

$$W_{\pi,h}^{\boldsymbol{\theta}}(s) = \mathbb{E}_{\pi} \left[ \sum_{h'=h}^{H} R(S_{h'}, A_{h'}) \middle| \boldsymbol{\theta} \right]$$

We define  $\pi^*(\theta)$  to be the optimal policy for some MDP (or contextual bandit) environment  $\theta$ ; barring computational issues, the optimal policy for a given MDP environment can be solved for using dynamic programming.

The frequentist regret is defined as the difference in total expected reward for the optimal policy versus the actual policy used:

$$\mathcal{R}_{K}(\pi,\theta) = \sum_{k=0}^{K} \sum_{s \in \mathcal{S}} \rho(s) \left( V_{\pi^{*},1}^{\theta}(s) - V_{\pi_{k},1}^{\theta}(s) \right)$$

Above  $\rho(s)$  represents the probability of starting the episode in state *s*, so  $\sum_{s \in S} \rho(s)$ . For the Bayesian regret, we assume the MDP environment  $\theta$  is drawn from prior distribution *Q*. The Bayesian regret is defined as follows:

$$\mathcal{BR}_K(\pi, \theta) = \mathbb{E}_{\theta \sim Q}[\mathcal{R}_K(\pi, \theta)]$$

Note, frequentist regret bounds are automatically Bayesian regret bounds, as they must hold for the worst case environment  $\theta$ . Bayesian regret bounds generally assume that the algorithm knows the prior on the environment Q.

**Theorem 1** (Bayesian Regret Bound for MDP Posterior Sampling). Let Q be the prior distribution over  $\theta$  used by the MDP posterior sampling algorithm. Let rewards  $R_t \in [0, C]$ , for some constant  $0 < C < \infty$ . Then for T = (K+1)H,

$$\mathcal{BR}_K(\pi,\theta) = \mathbb{E}_{\theta \sim Q}[\mathcal{R}_K(\pi,\theta)] = O(HS\sqrt{AT\log(SAT)})$$

Osband et al. [2013] prove that posterior sampling on MDPs has Bayesian regret  $\tilde{O}(HS\sqrt{AT})$ , as stated in Theorem 1 above. Since BHT-RL with posterior sampling contextual bandit and MDP algorithms is simply posterior sampling with a hierarchical prior, we can apply the regret bound of Theorem 1. Thus, BHT-RL with posterior sampling CB and MDP algorithms has Bayesian regret  $\tilde{O}(HS\sqrt{AT})$ , as stated in Corollary 1. In other words, Corollary 1 below follows directly from Theorem 1 because BHT-RL with CB and MDP algorithms is equivalent to posterior sampling with prior distribution  $Q := P(H_0)Q(\cdot|H_0) + P(H_1)Q(\cdot|H_1)$ .

**Corollary 1** (Bayesian Regret Bound for BHT-RL with Posterior Sampling). Suppose we use BHT-RL with posterior sampling CB and MDP algorithms. Let  $P(H_0) \in [0,1]$  be the prior probability of null hypothesis. Let  $Q(\cdot|H_0)$  and  $Q(\cdot|H_1)$  be the prior distribution over  $\theta$  conditional on the null and alternative hypotheses respectively. When rewards  $R_t \in [0, C]$ , for some constant  $0 < C < \infty$ ,

$$\mathcal{BR}_m(\pi,\theta) = \mathbb{E}_{\theta \sim O}[\mathcal{R}_m(\pi,\theta)] = O(HS\sqrt{AT\log(SAT)})$$

where distribution Q over  $\theta$  is defined as  $Q := P(H_0)Q(\cdot|H_0) + P(H_1)Q(\cdot|H_1).$ 

#### **4 EXPERIMENTS**

We run experiments in simulation environments to demonstrate the advantages of using BHT-RL when the nature of the environment is unknown. On the environments—a toy riverswim domain, randomly generated CBs / MDPs (see Appendix B), and a model inspired by a real world mobilehealth application—we show that BHT-RL significantly outperforms CB algorithms in MDP environments and vice versa, while performing nearly as well as CB algorithms in a bandit setting (and similarly for MDPs). Furthermore, we show our methods compare favorably with upper-bound based state of the arts methods aimed at adressing the same problem.

#### Figure 2: River Swim. Figure from Osband et al. [2013].

#### 4.1 ENVIRONMENTS

**Riverswim.** The first MDP environment we consider is the river swim environment introduced in Osband et al. [2013], and illustrated in Figure 2. The optimal policy in this environment must take a series of optimal actions to reach the high reward state on the right, and therefore a bandit algorithm which does not consider the long term benefits of actions will perform very poorly in it. In order to compare with a similar bandit environment, we construct a "CB River Swim environment" in which the the transition probability between any two states is uniform and independent of the action, while the rewards for each state are equivalent to those of the original MDP.

To test the performance of different algorithms as a function of their "banditness", we interpolate between the two environments by constructing domains with the following transition function:

$$P_{\lambda} = (1 - \lambda)P_{\rm CB} + \lambda P_{\rm MDP},\tag{1}$$

where  $P_{\text{CB}}$  and  $P_{\text{MDP}}$  are the transition functions for the CB and MDP environments respectively. Thus,  $\lambda = 1$  reduces to the original MDP environment, and as  $\lambda \rightarrow 0$ , the environment resembles more and more a CB, as the effect of actions on the transition probabilities diminishes.

Mobile Health Physical Activity Suggestions Environment. We consider a more realistic simulation environment, which is motivated by the mobile health problem of learning when to send activity suggestions to users. Highly sedentary lifestyles are associated with increased rates of many diseases including cardiovascular disease and diabetes [Hu, 2003, Biswas et al., 2015]. Health apps, through notifications delivered to mobile phones, smart watches and other wearable devices, have increasingly been used to remind users to take walks in order to encourage physical activity. RL is particularly important for learning when to send activity suggestions to users because of the high rate at which users stop using mobile health applications [Eysenbach, 2005]. The RL algorithm must learn to send messages when users will be receptive to activity suggestions and not send messages when users will find messages bothersome.

In this simulation setting, several times a day the RL algorithm must decide whether or not to send the user a message encouraging them to take a walk. The algorithm must learn in which contexts to send messages in order to maximize the physical activity of the user. The contextual information we include is detailed in Table 1. The reward is the log-step count in a fixed time period following the decision time. Our choice of features and reward is inspired by real world mobile health implementations, particularly Liao et al. [2020], who recently ran a mobile health study encouraging physical activity among people with hypertension. Additionally, the step count goal feature is inspired by the FitBit, which by default includes hourly step goals for users.



Figure 3: **Cumulative regret of BHT-RL compared with baselines.** The cumulative regret of the different algorithms is plotted vs. the number of episodes for the riverswim (top) and physical activity suggestions (bottom) environments. For each environment we plot the performance of its CB (left), MDP (right) and intermediate (center) variants. For the posterior sampling algorithms, the Dirichlet prior  $\alpha$  vector is set to all 1s. In all cases, BHT-RL's performance is comparable with the best performing algorithm. Error bars denote standard errors over 100 repetitions.

Context Variable	Values		
Time of Day	Morning, Afternoon, Evening		
Weather	Fair, Poor		
Engagement	Engaged, Disengaged		
Reached Step Goal	Goal Met, Goal Missed		

Table 1: Values contextual variables can take (24 total states).

In our simulation environment, the expected reward for sending a message is generally greater than or equal to the expected reward for not sending a message. However, while sending a message generally increases the immediate reward, it may increase the probability that the user transitions to a low reward state. In particular, if messages are sent when users are disengaged, this leads to a higher probability of transitioning to a disengaged state. However, if messages are sent when users are engaged, this leads to high reward and a high probability of remaining engaged. We use this model for the transition probabilities to reflect how users can both become more engaged over time when messages are sent at opportune times and how users can become more annoyed with messages if sent at inconvenient times. We construct two base environments-a CB and an MDP, and adjust the environment by modifying how much actions can affect state transition probabilities using the same method as in Equation (1), e.g., how much previous actions can affect users' future probability of being engaged with the app and



Figure 4: **Performance as a function of "banditness".** Cumulative regret at H = 100 relative to BHT-RL in the riverswim (left) and physical activity (right) environments vs. the interpolation parameter  $\lambda$  with Dirichlet prior  $\alpha = 1$ . Error bars denote standard errors over 100 repetitions.

reaching their step goals; see Appendix B.

In this environment we make use of the exogenous vs. endogenous decomposition of state describes in Section 3.3. Specifically, we treat time of day and weather as exogenous variables which are known to be unaffected by the agent's actions, and are therefore not included in the Bayesian hypothesis testing component.

#### 4.2 RESULTS

**BHT-RL consistently outperforms CB and MDP algorithms in their respective misspecified environments.** In Figure 9 we plot the cumulative regret over episodes for both environments. For each environments three variants



Figure 5: **Posterior probability of the null hypothesis.** The posterior probability of the null hypothesis in the riverswim (left) and physical activity (right) environments evolves with the number of episodes for different interpolation parameters  $\lambda$ . We set H = 100 and Dirichlet prior  $\alpha = 1$ . Error bars denote standard errors over 100 repetitions.



Figure 6: Maximum action variation for different environments over  $\lambda$ .

are used—the pure CB and pure MDP variants, as well as an intermediate environment in which the impact of actions on the transition probabilities is smaller than for the original MDP variant of the environment ( $\lambda = 0.6$ ). As noted in Figure 1, the performance of the CB and MDP posterior sampling (PS) algorithms is optimal when the algorithm is used for the correct environment (lower regret is better). However, we see that in all cases BHT-RL outperforms the CB and MDP posterior sampling algorithms in their respective misspecified environments, i.e.,  $\lambda$  close to 0 for MDP-PS and  $\lambda$  close to 1 for CB-PS. Moreover, in many cases the BHT-RL performs comparable to, if not better, than the algorithm explicitly designed for the particular environment. This demonstrates the ability of BHT-RL to perform well without knowledge of the nature of the environment. Furthermore, comparison with the UBEV-S [Zanette and Brunskill, 2018] and EULER [Zanette and Brunskill, 2019] baselines, which are also designed to operate when the nature of the environment is unknown, shows that they under-perform because of their reliance on confidence bounds which may be very loose.

We note that while the distinction between a bandit environment and an MDP is well defined, we can blur the definition and ask how "bandit-like" an MDP is by considering how strong of an effect do actions have on transition probabilities. In this sense, the  $\lambda$  parameter in equation (1) controls how close to a bandit our environment is, where the examples shown in Figure 1 correspond to  $\lambda = 0$  and 1. To better quantify the performance of the different posterior sampling methods as a function of the "banditness" of the environment, in Figure 4 we plot the cumulative regret after H = 100 episodes as a function of  $\lambda$ . For clarity of presentation, we plot the ratio between the regret for the CB or MDP algorithms over the regret of BHT-RL. BHT-RL consistently performs better than CB-PS and MDP-PS in their corresponding misspecified environments. For  $\lambda$  close to 1, BHT-RL outperforms CB-PS because through the Bayesian hypothesis testing procedure BHT-RL is able to learn that the environment is an MDP environment, while CB-PS is not able to ever learn the optimal policy. For  $\lambda$  close to 0, BHT-RL outperforms MDP-PS because the Bayesian hypothesis testing procedure regularizes the BHT-RL algorithm by reducing the planning horizon, which leads it to incur lower regret. BHT-RL consistently outperforms the worst performing algorithm (out of CB-PS and MDP-PS) across all values of  $\lambda$ . Moreover, BHT-RL has better or comparable performance to the best performing algorithm for most values of  $\lambda$ , with the exception of intermediate values of  $\lambda$  in the physical activity suggestion environment which we will be better equipped to discuss following the next section.

**Using BHT-RL to learn about the nature of the environment.** An additional benefit of BHT-RL is that it naturally outputs a posterior estimate for the probability of the null hypothesis that actions do not affect state transitions. In Figure 5 we plot the posterior probability of the null hypothesis as a function of the number of episodes. This knowledge can be useful in practice if we would like to consider what additional algorithms or methods to apply to the domain, and is obtained at no additional computational cost. This is in line with the Bayesian hypothesis testing literature (in contrast to frequentist hypothesis testing), which uses cutoff values for the Bayes factor to reject the null hypothesis [Quintana and Williams, 2018].

As a way to illustrate the relative difficulties of distinguishing between CB and MDP in the two simulation environments, following the terminology used in Jiang et al., we define the maximum action variation for an environment as the following:

$$\max_{a,a'\in\mathcal{A};s,s'\in\mathcal{S}}\left|P_{\theta}(s'|s,a)-P_{\theta}(s'|s,a')\right|$$

Note that for small non-zero values of  $\lambda$ , which correspond to smaller maximum action variation values, the posterior probability of  $H_0$  decays to zero very slowly, indicating that a large amount of data is needed to rule out  $H_0$  if the effect of actions on transition probabilities is small. As shown in Figure 6, the physical activity suggestions environment has a much smaller maximum action variation compared to river swim. The difficulty of learning whether the environment is a CB vs. MDP in the physical activity suggestions environment, is further demonstrated in the plots of the posterior



Figure 7: Sensitivity to algorithm hyper-parameters. We plot the cumulative regret at H = 100 as a function of the Dirichlet hyper-parameter  $\alpha$  for the riverswim (top) and physical activity suggestions (bottom) environments. For each environment we plot the results for its CB (left), MDP (right) and intermediate (center) variants.

probability of null hypothesis in Figure 5. We believe that the increased difficulty of learning whether the environment is a CB vs. MDP is what leads MDP-PS to outperform BHT-RL for intermediate values of  $\lambda$  in Figure 4.

Sensitivity to algorithm hyper-parameters. We demonstrate the sensitivity of our algorithm to  $\alpha$ . the parameter determining the Dirichlet distribution of the prior over transitions for the MDP algorithm. In Figure 7 we plot the regret of the different PS algorithms for the different environments. As expected, the prior on transition probabilities does not effect the performance in bandit settings, but may affect performance in an MDP environment where being able to model the transition probabilities well is important.

While the value of  $\alpha$  affects the performance of both the MDP-PS and BHT-RL algorithms, BHT-RL consistently performs better than or comparable to (1) CB algorithms in MDP environments and (2) MDP algorithms in CB environments for all values of  $\alpha$ . BHT-RL also performs at least as well as the best performing algorithm in each environment with the exception is for very small values of  $\alpha$ , where MDP-PS has lower cumulative regret. We believe this is due to the fact that a very small value of  $\alpha$  pushes the MDP-PS and the MDP component of BHT-RL to learn very sparse transition probabilities, which are inconsistent with the true environment. This in turn leads the BHT-RL algorithm to require more examples to learn whether the environment is an MDP or CB, and therefore incurs greater regret; see Figure 8 in the Appendix. Note that this can be mitigated by choosing a larger value of  $\alpha$ .

## 5 DISCUSSION

Our simulation results show that at least in finite state MDP and contextual bandit environments, the BHT-RL algorithm can perform well even when it is unknown whether the environment is that of an MDP or contextual bandit. Additionally, the BHT-RL approach allows practitioners to easily incorporate prior knowledge about the environment dynamics into their algorithm. Finally, since BHT-RL stochastically reduces the planning horizon, it can also be used as a regularization method for the full MDP based algorithm.

Some limitations of our work are that our method assumes the stationarity of the dynamics of the RL environment. Thus, our method is not robust to non-stationarity, which is often encountered in real world sequential decision making problems. To adapt the method to a continuous state setting would additionally require a model of the transition probabilities, which may be unrealistic to assume is known for real world problems. Additional open questions are whether it is possible to show additional theoretical guarantees regarding the BHT-RL algorithm, like a frequentist regret bound or a regret bound when the prior is misspecified.

Beyond just learning whether the environment is that of a contextual bandit or an MDP, we conjecture that bayesian hypothesis testing could also be used to address other aspects of reinforcement learning problems. One example is learning better state representations [Ortner et al., 2019], which is a major open problem in the reinforcement learning field [Dulac-Arnold et al., 2019].

#### References

- Dilip Arumugam, David Abel, Kavosh Asadi, Nakul Gopalan, Christopher Grimm, Jun Ki Lee, Lucas Lehnert, and Michael L Littman. Mitigating planner overfitting in model-based reinforcement learning. *arXiv preprint arXiv:1812.01129*, 2018.
- John Asmuth, Lihong Li, Michael L Littman, Ali Nouri, and David Wingate. A bayesian sampling approach to exploration in reinforcement learning. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 19–26. AUAI Press, 2009.
- Jim Berger. Bayesian hypothesis testing. 2012. URL https://cbms-mum.soe.ucsc.edu/ lecture2.pdf.
- Aviroop Biswas, Paul I Oh, Guy E Faulkner, Ravi R Bajaj, Michael A Silver, Marc S Mitchell, and David A Alter. Sedentary time and its association with risk for disease incidence, mortality, and hospitalization in adults: a systematic review and meta-analysis. *Annals of internal medicine*, 162(2):123–132, 2015.
- Rohan Chitnis and Tomás Lozano-Pérez. Learning compact models for planning with exogenous processes. In *Conference on Robot Learning*, pages 813–822. PMLR, 2020.
- Thomas Dietterich, George Trimponias, and Zhitang Chen. Discovering and removing exogenous state variables and rewards for reinforcement learning. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1262–1270, Stockholmsmässan, Stockholm Sweden, 10– 15 Jul 2018. PMLR. URL http://proceedings. mlr.press/v80/dietterich18a.html.
- Gabriel Dulac-Arnold, Daniel Mankowitz, and Todd Hester. Challenges of real-world reinforcement learning. *arXiv* preprint arXiv:1904.12901, 2019.
- Gunther Eysenbach. The law of attrition. *Journal of medical Internet research*, 7(1):e11, 2005.
- Frank B Hu. Sedentary lifestyle and risk of obesity and type 2 diabetes. *Lipids*, 38(2):103–108, 2003.
- Nan Jiang and Alekh Agarwal. Open problem: The dependence of sample complexity lower bounds on planning horizon. In *Conference On Learning Theory*, pages 3395– 3398, 2018.
- Nan Jiang, Satinder P Singh, and Ambuj Tewari. On structural properties of mdps that bound loss due to shallow planning.

- Nan Jiang, Alex Kulesza, Satinder Singh, and Richard Lewis. The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1181–1189. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E. Schapire. Contextual decision processes with low Bellman rank are PAC-learnable. In Doina Precup and Yee Whye Teh, editors, Proceedings of the 34th International Conference on Machine Learning, volume 70 of Proceedings of Machine Learning Research, pages 1704–1713, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL http://proceedings.mlr. press/v70/jiang17c.html.
- Peng Liao, Kristjan Greenewald, Predrag Klasnja, and Susan Murphy. Personalized heartsteps: A reinforcement learning algorithm for optimizing physical activity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(1):1–22, 2020.
- Ronald Ortner, Matteo Pirotta, Alessandro Lazaric, Ronan Fruit, and Odalric-Ambrym Maillard. Regret bounds for learning state representations in reinforcement learning. In Advances in Neural Information Processing Systems, pages 12717–12727, 2019.
- Ian Osband and Benjamin Van Roy. Why is posterior sampling better than optimism for reinforcement learning? In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2701–2710. JMLR. org, 2017.
- Ian Osband, Daniel Russo, and Benjamin Van Roy. (more) efficient reinforcement learning via posterior sampling. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, Advances in Neural Information Processing Systems 26, pages 3003-3011. Curran Associates, Inc., 2013. URL http://papers.nips.cc/paper/ 5185-more-efficient-reinforcement-learning-via pdf.
- Daniel S Quintana and Donald R Williams. Bayesian alternatives for common null-hypothesis significance tests in psychiatry: a non-technical guide using jasp. *BMC psychiatry*, 18(1):1–8, 2018.
- Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.
- Andrea Zanette and Emma Brunskill. Problem dependent reinforcement learning bounds which can identify bandit structure in MDPs. In Jennifer Dy and Andreas

Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5747–5755, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.

Andrea Zanette and Emma Brunskill. Tighter problemdependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, Proceedings of the 36th International Conference on Machine Learning, volume 97 of Proceedings of Machine Learning Research, pages 7304–7312, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL http://proceedings.mlr.press/v97/ zanette19a.html.

## A BAYESIAN HYPOTHESIS TESTING FOR DIRICHLET PRIORS ON TRANSITION PROBABILITIES

We define the set of states as S and the set of actions as A. Suppose we have data  $\mathcal{H}_T = \{S_t, A_t, R_t\}_{t=1}^T$ .

- Null hypothesis H<sub>0</sub>: We model our data as follows
  - For each  $s \in S$  we draw  $\tilde{\theta}_s \sim \text{Dirichlet}(\alpha)$
  - For all  $t \in [1: T]$  such that  $S_t = s$ , we have that  $S_{t+1} \sim \text{Categorical}(\tilde{\theta}_s)$
- Alternative hypothesis *H*<sub>1</sub>: We model our data as follows
  - For each  $s \in S$  and each  $a \in A$  we draw  $\tilde{\theta}_{s,a} \sim \text{Dirichlet}(\alpha)$
  - For all  $t \in [1: T]$  such that  $S_t = s$  and  $A_t = a$ , we have that  $S_{t+1} \sim \text{Categorical}(\tilde{\theta}_{s,a})$

We choose prior probabilities over the hypotheses  $P(H_0)$  and  $P(H_1) = 1 - P(H_0)$ . Then we can calculate the posterior probabilities  $P(H_0|\mathcal{H}_T)$  and  $P(H_1|\mathcal{H}_T)$ 

$$P(H_0|\mathcal{H}_T) = \frac{P(H_0, \mathcal{H}_T)}{P(\mathcal{H}_T)} = \frac{P(\mathcal{H}_T|H_0)P(H_0)}{P(\mathcal{H}_T|H_0)P(H_0) + P(\mathcal{H}_T|H_1)P(H_1)} = \frac{1}{1+K}$$

where  $K = \frac{P(\mathcal{H}_T | H_1) P(H_1)}{P(\mathcal{H}_T | H_0) P(H_0)}$  is the Bayes factor.

Let us now derive the posterior distributions. Let  $\theta$  represent all transition probability parameters, so  $\theta = {\{\theta_s\}_{s \in S} \cup {\{\theta_{s,a}\}_{s \in S, a \in A}}$ .

$$P(\theta|\mathcal{H}_T) = \frac{P(\theta, \mathcal{H}_T)}{P(\mathcal{H}_T)} = \frac{P(\mathcal{H}_T|\theta)P(\theta)}{\int P(\mathcal{H}_T|\theta)P(\theta)d\theta} =: \frac{X}{Y}$$

First examining the numerator term X,

$$X = P(\mathcal{H}_T|\theta) \left[ P(\theta|H_0)P(H_0) + P(\theta|H_1)P(H_1) \right]$$

$$= P(H_0) \prod_{s \in S} \left[ \text{Dirichlet}(\theta_s; \alpha) \prod_{t=1}^T \text{Categorical}(S_{t+1}; \theta_s)^{\mathbbm{1}_{S_{t}=s}} \right] \\ + P(H_1) \prod_{s \in S} \prod_{a \in A} \left[ \text{Dirichlet}(\theta_{s,a}; \alpha) \prod_{t=1}^T \text{Categorical}(S_{t+1}; \theta_{s,a})^{\mathbbm{1}_{S_t=s,A_t=a}} \right]$$

Below  $B(\alpha)$  is the multivariate beta function.

$$\begin{split} &= \frac{P(H_{0})}{B(\alpha)^{|S|}} \prod_{s \in S} \left[ \prod_{s' \in S} \theta_{s}(s')^{\alpha(s')-1} \prod_{t=1}^{T} \theta_{s}(s')^{\mathbbm{1}_{S_{t}=s,S_{t+1}=s'}} \right] \\ &\quad + \frac{P(H_{1})}{B(\alpha)^{|S||,|A|}} \prod_{s \in S} \prod_{a \in A} \left[ \prod_{s' \in S} \theta_{s,a}(s')^{\alpha(s')-1} \prod_{t=1}^{T} \theta_{s,a}(s')^{\mathbbm{1}_{S_{t}=s,A_{t}=a,S_{t+1}=s'}} \right] \\ &= \frac{P(H_{0})}{B(\alpha)^{|S|}} \prod_{s \in S} \left[ \prod_{s' \in S} \theta_{s}(s')^{\alpha(s')-1+\sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,S_{t+1}=s'}} \right] + \frac{P(H_{1})}{B(\alpha)^{|S||,|A|}} \prod_{s \in S} \prod_{a \in A} \left[ \prod_{s' \in S} \theta_{s,a}(s')^{\alpha(s')-1+\sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,A_{t}=a,S_{t+1}=s'}} \right] \\ &= \frac{P(H_{0})}{B(\alpha)^{|S|}} \prod_{s \in S} \left[ \prod_{s' \in S} \theta_{s}(s')^{\alpha(s')-1+\sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,S_{t+1}=s'}} \right] + \frac{P(H_{1})}{B(\alpha)^{|S||,|A|}} \prod_{s \in S} \prod_{a \in A} \left[ \prod_{s' \in S} \theta_{s,a}(s')^{\alpha(s')-1+\sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,A_{t}=a,S_{t+1}=s'}} \right] \\ &= \frac{P(H_{0})}{B(\alpha)^{|S|}} \prod_{s \in S} \left[ \prod_{s' \in S} \theta_{s}(s')^{\alpha(s')-1+\sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,S_{t+1}=s'}} \right] + \frac{P(H_{1})}{B(\alpha)^{|S||,|A|}} \prod_{s \in S} \prod_{a \in A} \left[ \prod_{s' \in S} \theta_{s,a}(s')^{\alpha(s')-1+\sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,A_{t}=a,S_{t+1}=s'}} \right] \\ &\text{We define } \mathbf{N}_{s} = \left[ \sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,S_{t+1}=1}, \sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,A_{t}=a,S_{t+1}=2}, \dots, \sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,A_{t}=a,S_{t+1}=|S|} \right] \text{ and } \\ \mathbf{N}_{s,a} = \left[ \sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,A_{t}=a,S_{t+1}=1}, \sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,A_{t}=a,S_{t+1}=2}, \dots, \sum_{t=1}^{T} \mathbbm{1}_{S_{t}=s,A_{t}=a,S_{t+1}=|S|} \right] \\ &= \frac{P(H_{0})}{B(\alpha)^{|S|}} \prod_{s \in S} B(\alpha + \mathbf{N}_{s}) \text{Dirichlet}(\theta_{s}; \alpha + \mathbf{N}_{s}) + \frac{P(H_{1})}{B(\alpha)^{|S||A|}} \prod_{s \in S} \prod_{a \in A} B(\alpha + \mathbf{N}_{s,a}) \text{Dirichlet}(\theta_{s,a}; \alpha + \mathbf{N}_{s,a}) \\ &= \frac{P(H_{0})}{B(\alpha)^{|S|}} \prod_{s \in S} B(\alpha + \mathbf{N}_{s}) \text{Dirichlet}(\theta_{s}; \alpha + \mathbf{N}_{s}) + \frac{P(H_{1})}{B(\alpha)^{|S||A|}} \prod_{s \in S} \prod_{a \in A} B(\alpha + \mathbf{N}_{s,a}) \text{Dirichlet}(\theta_{s,a}; \alpha + \mathbf{N}_{s,a}) \\ &= \frac{P(H_{0})}{B(\alpha)^{|S|}} \prod_{s \in S} m^{S} \prod_$$

Thus,

$$X = \frac{P(H_0)B(\alpha)^{|\mathcal{S}|(|\mathcal{A}|-1)}\prod_{s\in\mathcal{S}}B(\alpha+\mathbf{N}_s)\text{Dirichlet}(\theta_s;\alpha+\mathbf{N}_s)}{B(\alpha)^{|\mathcal{S}||\mathcal{A}|}} + \frac{P(H_1)\prod_{s\in\mathcal{S}}\prod_{a\in\mathcal{A}}B(\alpha+\mathbf{N}_{s,a})\text{Dirichlet}(\theta_{s,a};\alpha+\mathbf{N}_{s,a})}{B(\alpha)^{|\mathcal{S}||\mathcal{A}|}}$$

Since  $X = P(\mathcal{H}_T | \theta) P(\theta)$  and  $Y = \int P(\mathcal{H}_T | \theta) P(\theta) d\theta$ , we have that

$$Y = \frac{P(H_0)}{B(\alpha)^{|\mathcal{S}|}} \prod_{s \in \mathcal{S}} B(\alpha + \mathbf{N}_s) + \frac{P(H_1)}{B(\alpha)^{|\mathcal{S}||\mathcal{A}|}} \prod_{s \in \mathcal{S}} \prod_{a \in \mathcal{A}} B(\alpha + \mathbf{N}_{s,a})$$
$$= \frac{P(H_0)B(\alpha)^{|\mathcal{S}|(|\mathcal{A}|-1)} \prod_{s \in \mathcal{S}} B(\alpha + \mathbf{N}_s) + P(H_1) \prod_{s \in \mathcal{S}} \prod_{a \in \mathcal{A}} B(\alpha + \mathbf{N}_{s,a})}{B(\alpha)^{|\mathcal{S}||\mathcal{A}|}} =: \frac{W_0 + W_1}{B(\alpha)^{|\mathcal{S}||\mathcal{A}|}}$$

Thus,

$$P(\boldsymbol{\theta}|\mathcal{H}_T) = \frac{X}{Y} = \frac{W_0}{W_0 + W_1} \prod_{s \in S} \text{Dirichlet}(\boldsymbol{\theta}_s; \boldsymbol{\alpha} + \mathbf{N}_s) + \frac{W_0}{W_0 + W_1} \prod_{s \in S} \prod_{a \in A} \text{Dirichlet}(\boldsymbol{\theta}_{s,a}; \boldsymbol{\alpha} + \mathbf{N}_{s,a})$$

Note that

$$P(H_0|\mathcal{H}_T) = \frac{P(H_0|\mathcal{H}_T)P(\mathcal{H}_T)}{P(\mathcal{H}_T)} = \frac{P(\mathcal{H}_T|H_0)P(H_0)}{P(\mathcal{H}_T)} = \frac{W_0}{W_0 + W_1}$$
$$P(H_1|\mathcal{H}_T) = \frac{W_1}{W_0 + W_1}$$

## **B** SIMULATION DETAILS

## **B.1 RIVERSWIM ENVIRONMENTS**

- We add  $\mathcal{N}(0, 0.01)$  noise to all rewards.
- $P_{\text{MDP}}$  transitions are those of the original river swim environment as in Figure 2.
- $P_{\text{CB}}$  transitions are uniform over all states, i.e,  $P(S_{t+1} = s | S_t = s', A_t = a) = \frac{1}{|S|}$  for all  $s, s' \in S$  and  $a \in A$ .

## **Algorithm Hyper-Parameters**

- For Bandit and MDP posterior sampling we have independent  $\mathcal{N}(1,1)$  priors on the rewards.
- For MDP posterior sampling we have Dirichlet ( $\alpha \in \mathbb{R}^{S}$ ) priors on the transition probabilities.
- For BHT-PSRL we set the probability of the null hypothesis to  $P(H_0) = 0.5$ .
- UBEV-S and EULER we choose failure probability  $\delta = 0.1$ .

## **B.2 PHYSICAL ACTIVITY SUGGESTIONS ENVIRONMENTS**

- Reward is log step count
- Actions are binary: 0 means no message sent; 1 means message sent
- We add  $\mathcal{N}(0, 0.01)$  noise on rewards

Context Variable	Values	Variable Notation	Endogenous vs. Exogenous
Time of Day	Morning (0), Afternoon (1), Evening (2)	$S_t^{\text{time}}$	Exogenous
Weather	Fair (0), Poor (1)	$S_t^{\text{weather}}$	Exogenous
Engagement	Disengaged (0), Engaged (1)	$S_t^{\text{engagement}}$	Endogenous
Reached Step Goal	Goal Missed (0), Goal Met (1)	$S_t^{\text{goal}}$	Endogenous

Table 2: Values that contextual variables can take. There are 24 distinct states in total.

### Endogenous vs. Exogenous State Variables

• In our simulation environment, state space S can be decomposed into two parts:  $S = \mathcal{X} \times \mathcal{Z}$  such that  $\mathcal{X}$  state variables are exogenous and  $\mathcal{Z}$  state variables are potentially endogenous in the following sense:

$$P(S_{t+1}|S_t, A_t) = P(X_{t+1}, Z_{t+1}|X_t, Z_t, A_t) = P(X_{t+1}|X_t)P(Z_{t+1}|Z_t, A_t)$$

- For the purposes of Bayesian hypothesis PSRL approach, the bayesian hypothesis testing only has to occur for substates  $Z_t$ , rather than the full state  $S_t$ , which improves the performance and stability of our approach.
- For the model of the reward, we assume that  $E[R_t|S_t, A_t]$ , the expected reward in a given state and action, is affected by both  $X_t$  and  $Z_t$ . For example, how much a user walks following a decision time can be affected by both endogenous and exogenous variables.

### **State Transition Probabilities**

	Morning	Afternoon	Evening
Morning	0	1	0
Afternoon	0	0	1
Evening	1	0	0

Table 3: Time of Day Transition Probability Matrix

	Fair Weather	Poor Weather
Fair Weather	0.6	0.4
Poor Weather	0.3	0.7

Table 4: Weather Transition Probability Matrix

	Disengaged, Goal Missed	Disengaged, Goal Met	Engaged, Goal Missed	Engaged, Goal Met
Disengaged, Goal Missed	0.35	0.35	0.15	0.15
Disengaged, Goal Met	0.4	0.25	0.2	0.15
Engaged, Goal Missed	0.2	0.25	0.3	0.25
Engaged, Goal Met	0.15	0.15	0.3	0.4

Table 5: Endogenous Variable Transition Probability Matrix under  $A_t = 1$ , i.e.,  $P(Z_{t+1}|Z_t, A_t = 1)$ 

	Disengaged,	Disengaged,	Engaged,	Engaged,
Disongoood Cool Missod				
Diseligageu, Goal Misseu	0.43	0.33	0.1	0.1
Disengaged, Goal Met	0.5	0.3	0.15	0.05
Engaged, Goal Missed	0.05	0.3	0.3	0.35
Engaged, Goal Met	0.05	0.05	0.35	0.55

Table 6: Endogenous Variable Transition Probability Matrix under  $A_t = 0$ , i.e.,  $P(Z_{t+1}|Z_t, A_t = 0)$ 

- Above we state the transition probabilities under  $P_{\text{MDP}}$
- For transition probabilities  $P_{\text{CB}}$ , we simply set  $P(Z_{t+1}|Z_t, A_t = 1)$  to the corresponding values of  $P(Z_{t+1}|Z_t, A_t = 0)$ under P<sub>MDP</sub>.

#### **Expected Reward**

$$E[R_t|S_t, A_t = a] = \sum_{i=0}^{2} \theta_{a,i}^{\text{time}} \mathbb{1}_{i=S_t^{\text{time}}} + \sum_{j=0}^{1} \theta_{a,j}^{\text{weather}} \mathbb{1}_{j=S_t^{\text{weather}}} + \sum_{k=0}^{1} \sum_{l=0}^{1} \theta_{a,k,l}^{\text{endogenous}} \mathbb{1}_{k=S_t^{\text{engagement}}} \mathbb{1}_{l=S_t^{\text{goal}}}$$

• Time of Day

- 
$$\theta_{0,0}^{\text{time}} = 0.001; \ \theta_{0,1}^{\text{time}} = 0.01; \ \theta_{0,2}^{\text{time}} = 0.005$$

- 
$$\theta_{1,0}^{\text{time}} = 0.001; \ \theta_{1,1}^{\text{time}} = 0.02; \ \theta_{1,2}^{\text{time}} = 0.01$$

- Weather

  - $\begin{array}{l} \ \theta_{0,0}^{\text{weather}} = 0.01; \ \theta_{0,1}^{\text{weather}} = 0.015 \\ \ \theta_{1,0}^{\text{weather}} = 0.01; \ \theta_{1,1}^{\text{weather}} = 0.025 \end{array}$
- Endogenous

$$- \theta_{0,0,0}^{\text{endogenous}} = 0.005; \ \theta_{0,0,1}^{\text{endogenous}} = 0.4; \ \theta_{0,1,0}^{\text{endogenous}} = 0.35; \ \theta_{0,1,1}^{\text{endogenous}} = 2.25 \\ - \theta_{1,0,0}^{\text{endogenous}} = 0.01; \ \theta_{1,0,1}^{\text{endogenous}} = 0.405; \ \theta_{1,1,0}^{\text{endogenous}} = 1.75; \ \theta_{1,1,1}^{\text{endogenous}} = 2.5$$

#### **Algorithm Hyper-Parameters**

- For Bandit and MDP posterior sampling we have independent  $\mathcal{N}(1,1)$  priors on the rewards.
- For MDP posterior sampling we have Dirichlet ( $\alpha \in \mathbb{R}^{S}$ ) priors on the transition probabilities.
- For BHT-PSRL we set the probability of the null hypothesis to  $P(H_0) = 0.5$ .
- UBEV-S and EULER we choose failure probability  $\delta = 0.1$ .

## **B.3 RANDOM MDP ENVIRONMENTS**

- The following simulation environment is based on that in Jiang et al. [2015].
- For these experiments we randomly sampled 100 MDPs with 10 states and 2 actions from a distribution we refer to as Random-MDP, defined as follows.
- *P*<sub>MDP</sub> transitions are constructed as follows. For each *s* ∈ S and each *a* ∈ A, the distribution *P*(*s'*|*s*, *a*) for all *s'* ∈ S is chosen by selecting 5 non-zero entries uniformly from all 10 states, filling these 5 entries with values sampled from Uniform(0, 1), and then by normalizing the values to sum to 1.
- $P_{CB}$  transitions are constructed by modifying the  $P_{MDP}$  transition probabilities. In particular, for each  $s \in S$ , the transition probabilities P(s'|s, a = 1) is set to P(s'|s, a = 0) from the  $P_{MDP}$  transition probabilities.
- Start state is selected uniformly from all 10 states.
- Mean rewards were likewise sampled independently from Uniform(0, 1).
- Rewards have noise  $\mathcal{N}(0, 0.01)$ .
- Results averaged over 100 randomly sampled MDPs.

#### **Algorithm Hyper-Parameters**

- For Bandit and MDP posterior sampling we have independent  $\mathcal{N}(1,1)$  priors on the rewards.
- For MDP posterior sampling we have Dirichlet ( $\alpha \in \mathbb{R}^{S}$ ) priors on the transition probabilities.
- For BHT-PSRL we set the probability of the null hypothesis to  $P(H_0) = 0.5$ .
- UBEV-S and EULER we choose failure probability  $\delta = 0.1$ .

#### **B.4 ADDITIONAL SIMULATION RESULTS**



Figure 8: **Sensitivity to algorithm hyper-parameters.** We plot how the posterior probability of the null hypothesis for different values of the Dirichlet hyper-parameter  $\alpha$  for the riverswim (Top), physical activity suggestions (Middle), and random MDP (Bottom) environments. For each environment we plot the results for its CB (Left), MDP (Right) and intermediate (Center) variants.



Figure 9: **Cumulative regret of BHT-RL compared with baselines.** The cumulative regret of the different algorithms is plotted vs. the number of episodes for the riverswim (Top) and physical activity suggestions (Bottom) environments. For each environment we plot the performance of its CB (Left), MDP (Right) and intermediate (Center) variants. For the posterior sampling algorithms, the Dirichlet prior  $\alpha$  vector is set to all 1s. In all cases, BHT-RL's performance is comparable with the best performing algorithm.



Figure 10: We plot regret relative to PSRL-HT in random MDP environments varying interpolation parameter  $\lambda$  with H = 100. Error bars denote standard errors over 100 repetitions. Dirichlet prior hyperparameter  $\alpha = 1$ .



Figure 11: We plot posterior probabilities for PSRL-HT in the random MDP environments for different interpolation parameters  $\lambda$  with H = 100. Error bars denote standard errors over 100 repetitions. Dirichlet prior hyperparameter  $\alpha = 1$ .



Figure 12: Left: Random CB environments. Middle: transition probabilities interpolated between Random MDP and Random CB environments; Right: Random MDP environments. All environments have horizon H = 100. Error bars denote standard errors over 100 repetitions. We vary dirichlet prior hyperparameter  $\alpha$ , letting it take values 0.25, 0.5, 1, 2, 4, 8.