
Deciding What to Model: Value-Equivalent Sampling for Reinforcement Learning

Dilip Arumugam

Department of Computer Science
Stanford University
dilip@cs.stanford.edu

Benjamin Van Roy

Department of Electrical Engineering
Department of Management Science & Engineering
Stanford University
bvr@stanford.edu

Abstract

The quintessential model-based reinforcement-learning agent iteratively refines its estimates or prior beliefs about the true underlying model of the environment. Recent empirical successes in model-based reinforcement learning with function approximation, however, eschew the true model in favor of a surrogate that, while ignoring various facets of the environment, still facilitates effective planning over behaviors. Recently formalized as the value equivalence principle, this algorithmic technique is perhaps unavoidable as real-world reinforcement learning demands consideration of a simple, computationally-bounded agent interacting with an overwhelmingly complex environment, whose underlying dynamics likely exceed the agent’s capacity for representation. In this work, we consider the scenario where agent limitations may entirely preclude identifying an exactly value-equivalent model, immediately giving rise to a trade-off between identifying a model that is simple enough to learn while only incurring bounded sub-optimality. To address this problem, we introduce an algorithm that, using rate-distortion theory, iteratively computes an approximately-value-equivalent, lossy compression of the environment which an agent may feasibly target in lieu of the true model. We prove an information-theoretic, Bayesian regret bound for our algorithm that holds for any finite-horizon, episodic sequential decision-making problem. Crucially, our regret bound can be expressed in one of two possible forms, providing a performance guarantee for finding either the simplest model that achieves a desired sub-optimality gap or, alternatively, the best model given a limit on agent capacity.

1 Introduction

A central challenge of the reinforcement-learning problem [154, 87] is exploration, where a sequential decision-making agent must judiciously balance exploitation of knowledge accumulated thus far against the need to further acquire information for optimal long-term performance. Historically, provably-efficient reinforcement-learning algorithms [91, 35, 90, 20, 23, 151, 81, 121, 45, 120, 22, 46, 10, 85, 168, 56, 105] have often relied upon one of two possible mechanisms for addressing the exploration challenge in a principled manner: optimism in the face of uncertainty or posterior sampling. Briefly, methods in the former category begin with optimistically-biased value estimates for all state-action pairs; an agent acting greedily with respect to these estimates will be incentivized to visit all state-action pairs a sufficient number of times until this bias dissipates and the agent is left with an accurate estimate of the value function for deriving optimal behavior. In contrast, posterior-sampling methods primarily operate based on Thompson sampling [156, 141] whereby the agent begins with a prior belief over the Markov Decision Process (MDP) with which it is interacting and acts optimally with respect to a single sample drawn from these beliefs. The resulting experience sampled from the true environment allows the agent to derive a corresponding posterior distribution

and this Posterior Sampling for Reinforcement Learning (PSRL) [152] algorithm proceeds iteratively in this manner, eventually arriving at a posterior sharply concentrated around the true environment MDP. While both paradigms have laid down solid theoretical foundations for provably-efficient reinforcement learning, a line of work has demonstrated how posterior-sampling methods can be more favorable both in theory and in practice [121–123, 120, 114, 125, 63].

While existing analyses of reinforcement-learning algorithms have largely focused on providing guarantees for learning optimal solutions, real-world reinforcement learning demands consideration for a computationally-bounded agent interacting with an overwhelmingly complex environment [105]. A simplified view of this notion can be succinctly depicted in the multi-armed bandit setting [97, 38, 100]; as the number of arms increases, a Thompson sampling agent’s relentless pursuit of the optimal arm will lead to large regret [138, 139]. On the other hand, one might simply settle for the first ε -optimal arm found, for some $\varepsilon > 0$, which may be identified in far fewer time periods. The goal of this work is to augment PSRL so as to accommodate these satisficing solutions in addition to optimal ones, paralleling existing work for satisficing in multi-armed bandit problems [140, 139, 15, 16]. To help elucidate the utility of satisficing solutions in the reinforcement-learning setting, we offer the following illustrative example:

Example 1 (A Multi-Resolution MDP). *For a large but finite $N \in \mathbb{N}$, consider a sequence of MDPs, $\{\mathcal{M}_n\}_{n \in [N]}$, which all share a common action space \mathcal{A} but vary in state space \mathcal{S}_n , reward function, and transition function. Moreover, for each $n \in [N]$, the rewards of the n th MDP are bounded in the interval $[0, \frac{1}{n}]$. An agent is confronted with the resulting product MDP, \mathcal{M} , defined on the state space $\mathcal{S}_1 \times \dots \times \mathcal{S}_N$ with action space \mathcal{A} and rewards summed across the N constituent reward functions. The transition function is defined such that each action $a \in \mathcal{A}$ is executed across all N MDPs simultaneously and the resulting individual transitions are combined into a transition of \mathcal{M} .*

Example 1 presents a simple scenario where, as $N \uparrow \infty$, a complex environment retains a wealth of information and yet, due to the scale of N and the boundedness of rewards for each constituent MDP \mathcal{M}_n , only a subset of that information is within the agent’s reach or even necessary for producing reasonably competent behavior. Despite this fact, PSRL will persistently act to fully identify the transition and reward structure of all $\{\mathcal{M}_n\}_{n \in [N]}$, for any value of N . Without knowing which MDPs are more important *a priori* and even as data accumulates during learning, PSRL is unable to forego learning granular components of \mathcal{M} , eventually accumulating optimal reward at the cost of more time. Intuitively, however, one might anticipate that there exists a value $M \ll N$ such that learning the subsequence of MDPs $\{\mathcal{M}_n\}_{n \in [M]}$ in fewer time periods is sufficient for achieving a desired degree of sub-optimality, since the rewards of the remaining MDPs $\{\mathcal{M}_n\}_{n > M}$ make suitably negligible contributions to the overall rewards of \mathcal{M} . Alternatively, for a computationally-bounded decision maker, the agent’s resource limitations ought to translate into a value $C \ll N$ such that $\{\mathcal{M}_n\}_{n \in [C]}$ is feasible and learning this subsequence is the best possible outcome under the agent capacity constraints. In this work, we introduce an algorithm that, in a purely data-driven and automated fashion, implicitly identifies such a value M or C to facilitate tractable, near-optimal learning in what may otherwise be an intractable problem. Following Arumugam and Van Roy [15], a key tool for defining a notion of satisficing in reinforcement learning will be rate-distortion theory [146, 25].

The paper proceeds as follows: we introduce our problem formulation in Section 3, present our generalization of PSRL in Section 4, and provide a complementary regret analysis in Section 5. Due to space constraints, technical proofs, an overview of related work, and discussion of our results in a broader context are relegated to the appendix.

2 Preliminaries

In this section, we provide brief background on information theory and details on our notation. All random variables are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. For any random variable $X : \Omega \rightarrow \mathcal{X}$ taking values on the measurable space $(\mathcal{X}, \mathbb{X})$, we use $\sigma(X) \triangleq \{X^{-1}(A) \mid A \in \mathbb{X}\} \subseteq \mathcal{F}$ to denote the σ -algebra generated by X . For any natural number $N \in \mathbb{N}$, we denote the index set as $[N] \triangleq \{1, 2, \dots, N\}$. For any arbitrary set \mathcal{X} , $\Delta(\mathcal{X})$ denotes the set of all probability distributions with support on \mathcal{X} . For any two arbitrary sets \mathcal{X} and \mathcal{Y} , we denote the class of all (measurable) functions mapping from \mathcal{X} to \mathcal{Y} as $\{\mathcal{X} \rightarrow \mathcal{Y}\} \triangleq \{f \mid f : \mathcal{X} \rightarrow \mathcal{Y}\}$. While our exposition throughout the paper will consistently refer to bits of information, it will be useful for the purposes of analysis that all logarithms be in base e .

2.1 Information Theory

Here we introduce various concepts in probability theory and information theory used throughout this paper. We encourage readers to consult [41, 71, 129, 59] for more background.

We define the mutual information between any two random variables X, Y through the Kullback-Leibler (KL) divergence:

$$\mathbb{I}(X; Y) = D_{\text{KL}}(\mathbb{P}((X, Y) \in \cdot) \parallel \mathbb{P}(X \in \cdot) \times \mathbb{P}(Y \in \cdot)) \quad D_{\text{KL}}(P \parallel Q) = \begin{cases} \int \log \left(\frac{dP}{dQ} \right) dP & P \ll Q \\ +\infty & P \not\ll Q \end{cases},$$

where P and Q are both probability measures on the same measurable space and $\frac{dP}{dQ}$ denotes the Radon-Nikodym derivative of P with respect to Q . An analogous definition of conditional mutual information holds through the expected KL-divergence for any three random variables X, Y, Z :

$$\mathbb{I}(X; Y \mid Z) = \mathbb{E} [D_{\text{KL}}(\mathbb{P}((X, Y) \in \cdot \mid Z) \parallel \mathbb{P}(X \in \cdot \mid Z) \times \mathbb{P}(Y \in \cdot \mid Z))].$$

With these definitions in hand, we may define the entropy and conditional entropy for any two random variables X, Y as

$$\mathbb{H}(X) = \mathbb{I}(X; X) \quad \mathbb{H}(Y \mid X) = \mathbb{H}(Y) - \mathbb{I}(X; Y).$$

This yields the following identities for mutual information and conditional mutual information for any three arbitrary random variables X, Y , and Z :

$$\mathbb{I}(X; Y) = \mathbb{H}(X) - \mathbb{H}(X \mid Y) = \mathbb{H}(Y) - \mathbb{H}(Y \mid X), \quad \mathbb{I}(X; Y \mid Z) = \mathbb{H}(X \mid Z) - \mathbb{H}(X \mid Y, Z) = \mathbb{H}(Y \mid Z) - \mathbb{H}(Y \mid X, Z).$$

Through the chain rule of the KL-divergence and the fact that $D_{\text{KL}}(P \parallel P) = 0$ for any probability measure P , we obtain another equivalent definition of mutual information,

$$\mathbb{I}(X; Y) = \mathbb{E} [D_{\text{KL}}(\mathbb{P}(Y \in \cdot \mid X) \parallel \mathbb{P}(Y \in \cdot))],$$

as well as the chain rule of mutual information: $\mathbb{I}(X; Y_1, \dots, Y_n) = \sum_{i=1}^n \mathbb{I}(X; Y_i \mid Y_1, \dots, Y_{i-1})$.

Finally, for any three random variables X, Y , and Z which form the Markov chain $X \rightarrow Y \rightarrow Z$, we have the following data-processing inequality: $\mathbb{I}(X; Z) \leq \mathbb{I}(X; Y)$.

3 Problem Formulation

We formulate a sequential decision-making problem as a finite-horizon, episodic Markov Decision Process (MDP) [24, 130] defined by $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \beta, H \rangle$. Here \mathcal{S} denotes a set of states, \mathcal{A} is a set of actions, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is a deterministic reward function providing evaluative feedback signals (in the unit interval) to the agent, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is a transition function prescribing distributions over next states, $\beta \in \Delta(\mathcal{S})$ is an initial state distribution, and $H \in \mathbb{N}$ is the maximum episode length or horizon.

As is standard in Bayesian reinforcement learning [70], neither the transition function nor the reward function are known to the agent and, consequently, both are treated as random variables. Since all other components of the MDP are thought of as known a priori, the randomness in the model $(\mathcal{R}, \mathcal{T})$ fully accounts for the randomness in \mathcal{M} , which is also a random variable. We denote by \mathcal{M}^* the true MDP with model $(\mathcal{R}^*, \mathcal{T}^*)$ that the agent interacts with and attempts to solve over the course of K episodes. Within each episode, the agent acts for exactly H steps beginning with an initial state $s_1 \sim \beta$. For each $h \in [H]$, the agent observes the current state $s_h \in \mathcal{S}$, selects action $a_h \sim \pi_h(\cdot \mid s_h) \in \mathcal{A}$, enjoys a reward $r_h = \mathcal{R}(s_h, a_h) \in [0, 1]$, and transitions to the next state $s_{h+1} \sim \mathcal{T}(\cdot \mid s_h, a_h) \in \mathcal{S}$.

A stationary, stochastic policy for timestep $h \in [H]$, $\pi_h : \mathcal{S} \rightarrow \Delta(\mathcal{A})$, encodes a pattern of behavior mapping individual states to distributions over possible actions. Letting $\{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}$ denote the class of all stationary, stochastic policies, a non-stationary policy $\pi = (\pi_1, \dots, \pi_H) \in \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}^H$ is a collection of exactly H stationary, stochastic policies whose overall performance in any MDP \mathcal{M} at timestep $h \in [H]$ when starting at state $s \in \mathcal{S}$ and taking action $a \in \mathcal{A}$ is assessed by its

associated action-value function $Q_{\mathcal{M}, h}^\pi(s, a) = \mathbb{E} \left[\sum_{h'=h}^H \mathcal{R}(s_{h'}, a_{h'}) \mid s_h = s, a_h = a \right]$, where the

expectation integrates over randomness in the action selections and transition dynamics. Taking the corresponding value function as $V_{\mathcal{M},h}^\pi(s) = \mathbb{E}_{a \sim \pi_h(\cdot|s)} \left[Q_{\mathcal{M},h}^\pi(s, a) \right]$, we define the optimal policy $\pi^* = (\pi_1^*, \pi_2^*, \dots, \pi_H^*)$ as achieving supremal value $V_{\mathcal{M},h}^*(s) = \sup_{\pi \in \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}^H} V_{\mathcal{M},h}^\pi(s)$ for all $s \in \mathcal{S}$, $h \in [H]$. For brevity, we will write any value function $V \in \{\mathcal{S} \rightarrow \mathbb{R}\}$ without its argument to implicitly integrate over randomness in the initial state: $V = \mathbb{E}_{s_1 \sim \beta(\cdot)} [V(s_1)]$. We let $\tau_k = (s_1^{(k)}, a_1^{(k)}, r_1^{(k)}, \dots, s_H^{(k)}, a_H^{(k)}, r_H^{(k)}, s_{H+1}^{(k)})$ be the random variable denoting the trajectory experienced by the agent in the k th episode. Meanwhile, $H_k = \{\tau_1, \tau_2, \dots, \tau_{k-1}\} \in \mathcal{H}_k$ is the random variable representing the entire history of the agent's interaction within the environment at the start of the k th episode; the sequence of history random variables $\{H_k\}_{k \in [K]}$ induce and, by definition, are adapted to the filtration $\{\sigma(H_k)\}_{k \in [K]}$ of (Ω, \mathcal{F}) . We call attention to the fact that we have yet to make any further restrictions on the state-action space $\mathcal{S} \times \mathcal{A}$, such as finiteness; notably, the main results of this paper are not limited to tabular MDPs. As mentioned by Lattimore and Szepesvári [100] (also as Proposition 7.28 of Bertsekas and Shreve [28]), the Ionescu-Tulcea Theorem [80] ensures the existence of a probability space upon which τ_k and H_k are well-defined random variables for all episodes $k \in [K]$.

Abstractly, a reinforcement-learning algorithm is a sequence of non-stationary policies $(\pi^{(1)}, \dots, \pi^{(K)})$ where for each episode $k \in [K]$, $\pi^{(k)} : \mathcal{H}_k \rightarrow \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}^H$ is a function of the current history H_k . We define the regret of a reinforcement-learning algorithm over K episodes as

$$\text{REGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}, \mathcal{M}^*) = \sum_{k=1}^K \Delta_k \quad \Delta_k \triangleq V_{\mathcal{M}^*,1}^* - V_{\mathcal{M}^*,1}^{\pi^{(k)}}$$

where Δ_k denotes the episodic regret or regret incurred during the k th episode with respect to the true MDP \mathcal{M}^* . An agent's initial uncertainty in the (unknown) true MDP \mathcal{M}^* is reflected by an arbitrary prior distribution $\mathbb{P}(\mathcal{M}^* \in \cdot | H_1)$. Since the regret is a random variable due to our uncertainty in \mathcal{M}^* , we integrate over this randomness to arrive at the Bayesian regret:

$$\text{BAYESREGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}) = \mathbb{E} \left[\text{REGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}, \mathcal{M}^*) \right].$$

Broadly speaking, our goal is to design a provably-efficient reinforcement-learning algorithm that incurs bounded Bayesian regret.

Throughout the paper, we will denote the entropy and conditional entropy conditioned upon a specific realization of an agent's history H_k , for some episode $k \in [K]$, as $\mathbb{H}_k(X) \triangleq \mathbb{H}(X | H_k = H_k)$ and $\mathbb{H}_k(X | Y) \triangleq \mathbb{H}_k(X | Y, H_k = H_k)$, for two arbitrary random variables X and Y . This notation will also apply analogously to the mutual information $\mathbb{I}_k(X; Y) \triangleq \mathbb{I}(X; Y | H_k = H_k) = \mathbb{H}_k(X) - \mathbb{H}_k(X | Y) = \mathbb{H}_k(Y) - \mathbb{H}_k(Y | X)$, as well as the conditional mutual information $\mathbb{I}_k(X; Y | Z) \triangleq \mathbb{I}(X; Y | H_k = H_k, Z)$, given an arbitrary third random variable, Z . Note that their dependence on the realization of random history H_k makes both $\mathbb{I}_k(X; Y)$ and $\mathbb{I}_k(X; Y | Z)$ random variables themselves. The traditional notion of conditional mutual information given the random variable H_k arises by integrating over this randomness:

$$\mathbb{E} [\mathbb{I}_k(X; Y)] = \mathbb{I}(X; Y | H_k) \quad \mathbb{E} [\mathbb{I}_k(X; Y | Z)] = \mathbb{I}(X; Y | H_k, Z).$$

Additionally, we will also adopt a similar notation to express a conditional expectation given the random history H_k : $\mathbb{E}_k[X] \triangleq \mathbb{E}[X | H_k]$.

4 Satisficing Through Posterior Sampling

4.1 Rate-Distortion Theory

We begin with a brief, high-level overview of rate-distortion theory [146, 25] and encourage readers to consult [41] for more details and [26] for a survey of advances in rate-distortion theory towards solving the lossy source coding problem in information theory. A lossy compression problem consumes as input a fixed information source $\mathbb{P}(X \in \cdot)$ and a measurable distortion function $d : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}_{\geq 0}$ which quantifies the loss of fidelity by using Z in place of X . Then, for any

$D \in \mathbb{R}_{\geq 0}$, the rate-distortion function quantifies the fundamental limit of lossy compression as

$$\mathcal{R}(D) = \inf_{Z \in \Lambda(D)} \mathbb{I}(X; Z) \quad \Lambda(D) \triangleq \{Z : \Omega \rightarrow \mathcal{Z} \mid \mathbb{E}[d(X, Z)] \leq D\},$$

where the infimum is taken over all random variables Z that incur bounded expected distortion, $\mathbb{E}[d(X, Z)] \leq D$. Naturally, $\mathcal{R}(D)$ represents the minimum number of bits of information that must be retained from X in order to achieve this bounded expected loss of fidelity¹. Throughout the paper, various facts of the rate-distortion function will be referenced as needed. For now, we simply note that, in keeping with the problem formulation of the previous section which does not automatically assume discrete random variables, the rate-distortion function is well-defined for abstract information source and channel output random variables [43].

Just as in past work that studies satisficing in multi-armed bandit problems [138, 139, 15], we will use rate-distortion theory to formalize and identify the best simplified MDP $\widehat{\mathcal{M}}_k$ that the agent will attempt to learn over the course of each episode $k \in [K]$. The dependence on the particular episode comes from the fact that this lossy compression mechanism or channel will treat the agent's current beliefs over the true MDP $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_k)$ as the information source to be compressed.

4.2 The Value Equivalence Principle

As outlined in the previous section, the second input for a well-specified lossy-compression problem is a distortion function prescribing non-negative real values to realizations of the information source and channel output random variables $(\mathcal{M}^*, \widehat{\mathcal{M}})$ that quantify the loss of fidelity incurred by using $\widehat{\mathcal{M}}$ in lieu of \mathcal{M}^* . To define this function, we will leverage an approximate notion of value equivalence [72, 73]. For any arbitrary MDP \mathcal{M} with model $(\mathcal{R}, \mathcal{T})$ and any stationary, stochastic policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$, define the Bellman operator $\mathcal{B}_{\mathcal{M}}^{\pi} : \{\mathcal{S} \rightarrow \mathbb{R}\} \rightarrow \{\mathcal{S} \rightarrow \mathbb{R}\}$ as follows:

$$\mathcal{B}_{\mathcal{M}}^{\pi} V(s) \triangleq \mathbb{E}_{a \sim \pi(\cdot \mid s)} [\mathcal{R}(s, a) + \mathbb{E}_{s' \sim \mathcal{T}(\cdot \mid s, a)} [V(s')]], \quad \forall s \in \mathcal{S}.$$

The Bellman operator is a foundational tool in dynamic-programming approaches to reinforcement learning [29] and gives rise to the classic Bellman equation: for any MDP $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \beta, H \rangle$ and any non-stationary policy $\pi = (\pi_1, \dots, \pi_H)$, the value functions induced by π satisfy $V_{\mathcal{M}, h}^{\pi}(s) = \mathcal{B}_{\mathcal{M}}^{\pi_h} V_{\mathcal{M}, h+1}^{\pi}(s)$, for all $h \in [H]$ and with $V_{\mathcal{M}, H+1}^{\pi}(s) = 0, \forall s \in \mathcal{S}$. For any two MDPs $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \beta, H \rangle$ and $\widehat{\mathcal{M}} = \langle \mathcal{S}, \mathcal{A}, \widehat{\mathcal{R}}, \widehat{\mathcal{T}}, \beta, H \rangle$, Grimm et al. [72] define a notion of equivalence between them despite their differing models. For any policy class $\Pi \subseteq \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}$ and value function class $\mathcal{V} \subseteq \{\mathcal{S} \rightarrow \mathbb{R}\}$, \mathcal{M} and $\widehat{\mathcal{M}}$ are value equivalent with respect to Π and \mathcal{V} if and only if $\mathcal{B}_{\mathcal{M}}^{\pi} V = \mathcal{B}_{\widehat{\mathcal{M}}}^{\pi} V, \forall \pi \in \Pi, V \in \mathcal{V}$. In words, two different models are deemed value equivalent if they induce identical Bellman updates under any pair of policy and value function from $\Pi \times \mathcal{V}$. Grimm et al. [72] prove that when $\Pi = \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}$ and $\mathcal{V} = \{\mathcal{S} \rightarrow \mathbb{R}\}$, the set of all exactly value-equivalent models is a singleton set containing only the true model of the environment. The key insight behind value equivalence, however, is that practical model-based reinforcement-learning algorithms need not be concerned with modeling every granular detail of the underlying environment and may, in fact, stand to benefit by optimizing an alternative criterion besides the traditional maximum-likelihood objective [147, 66, 116, 17, 65, 57, 1, 44, 21, 144, 110, 113, 161]. Indeed, by restricting focus to decreasing subsets of policies $\Pi \subset \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}$ and value functions $\mathcal{V} \subset \{\mathcal{S} \rightarrow \mathbb{R}\}$, the space of exactly value-equivalent models is monotonically increasing.

For brevity, let $\mathfrak{R} \triangleq \{\mathcal{S} \times \mathcal{A} \rightarrow [0, 1]\}$ and $\mathfrak{T} \triangleq \{\mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})\}$ denote the classes of all reward functions and transition functions, respectively. Recall that, with $(\mathcal{S}, \mathcal{A}, \beta, H)$ all known, the uncertainty in a random MDP \mathcal{M} is entirely driven by its model $(\mathcal{R}, \mathcal{T})$ such that we may think of the support of \mathcal{M}^* as $\text{supp}(\mathcal{M}^*) = \mathfrak{M} \triangleq \mathfrak{R} \times \mathfrak{T}$. We define a distortion function on pairs of MDPs $d : \mathfrak{M} \times \mathfrak{M} \rightarrow \mathbb{R}_{\geq 0}$ for any $\Pi \subseteq \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}, \mathcal{V} \subseteq \{\mathcal{S} \rightarrow \mathbb{R}\}$ as

$$d_{\Pi, \mathcal{V}}(\mathcal{M}, \widehat{\mathcal{M}}) = \sup_{\substack{\pi \in \Pi \\ V \in \mathcal{V}}} \|\mathcal{B}_{\mathcal{M}}^{\pi} V - \mathcal{B}_{\widehat{\mathcal{M}}}^{\pi} V\|_{\infty}^2 = \sup_{\substack{\pi \in \Pi \\ V \in \mathcal{V}}} \left(\sup_{s \in \mathcal{S}} |\mathcal{B}_{\mathcal{M}}^{\pi} V(s) - \mathcal{B}_{\widehat{\mathcal{M}}}^{\pi} V(s)| \right)^2.$$

In words, $d_{\Pi, \mathcal{V}}$ is the supremal squared Bellman error between MDPs \mathcal{M} and $\widehat{\mathcal{M}}$ across all states $s \in \mathcal{S}$ with respect to the policy class Π and value function class \mathcal{V} .

¹With a slight abuse of notation, we overload \mathcal{R} .

4.3 Value-Equivalent Sampling for Reinforcement Learning

By virtue of the previous two sections, we are now in a position to define the lossy compression problem that characterizes a MDP $\widetilde{\mathcal{M}}$ that the agent will aspire to learn in each episode $k \in [K]$ instead of the true MDP \mathcal{M}^* . For any $\Pi \subseteq \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}$; $\mathcal{V} \subseteq \{\mathcal{S} \rightarrow \mathbb{R}\}$; $k \in [K]$; and $D \geq 0$, we define the rate-distortion function

$$\mathcal{R}_k^{\Pi, \mathcal{V}}(D) = \inf_{\widetilde{\mathcal{M}} \in \Lambda_k(D)} \mathbb{I}_k(\mathcal{M}^*; \widetilde{\mathcal{M}}), \quad \Lambda_k(D) \triangleq \{\widetilde{\mathcal{M}} : \Omega \rightarrow \mathfrak{M} \mid \mathbb{E}_k[d_{\Pi, \mathcal{V}}(\mathcal{M}^*, \widetilde{\mathcal{M}})] \leq D\}. \quad (1)$$

This rate-distortion function characterizes the fundamental limit of MDP compression under our chosen distortion measure resulting in a channel that retains the minimum amount of information from the true MDP \mathcal{M}^* while yielding an approximately value-equivalent MDP in expectation. Observe that this distortion constraint is a notion of approximate value equivalence which collapses to the exact value equivalence of Grimm et al. [72] as $D \rightarrow 0$. Meanwhile, as $D \rightarrow \infty$, we accommodate a more aggressive compression of the true MDP \mathcal{M}^* resulting in less faithful Bellman updates.

Algorithm 1 Posterior Sampling for Reinforcement Learning (PSRL) [152]

Input: Prior $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_1)$
for $k \in [K]$ **do**
 Sample $M_k \sim \mathbb{P}(\mathcal{M}^* \in \cdot \mid H_k)$
 Get optimal policy $\pi^{(k)} = \pi_{M_k}^*$
 Execute $\pi^{(k)}$ and get trajectory τ_k
 Update history $H_{k+1} = H_k \cup \tau_k$
 Induce posterior $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_{k+1})$
end for

Algorithm 2 Value-equivalent Sampling for Reinforcement Learning (VSRL)

Input: Prior $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_1)$, Threshold $D \in \mathbb{R}_{\geq 0}$, Distortion function $d_{\Pi, \mathcal{V}} : \mathfrak{M} \times \mathfrak{M} \rightarrow \mathbb{R}_{\geq 0}$
for $k \in [K]$ **do**
 Compute $\widetilde{\mathcal{M}}_k$ achieving $\mathcal{R}_k^{\Pi, \mathcal{V}}(D)$ limit (Equation 1)
 Sample MDP $M^* \sim \mathbb{P}(\mathcal{M}^* \in \cdot \mid H_k)$
 Sample compression $M_k \sim \mathbb{P}(\widetilde{\mathcal{M}}_k \in \cdot \mid \mathcal{M}^* = M^*)$
 Compute optimal policy $\pi^{(k)} = \pi_{M_k}^*$
 Execute $\pi^{(k)}$ and observe trajectory τ_k
 Update history $H_{k+1} = H_k \cup \tau_k$
 Induce posterior $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_{k+1})$
end for

A standard algorithm for our problem setting is widely known as Posterior Sampling for Reinforcement Learning (PSRL) [152, 120], which we present as Algorithm 1, while our Value-equivalent Sampling for Reinforcement Learning (VSRL) is given as Algorithm 2. The key distinction between them is that, at each episode $k \in [K]$, the latter takes the posterior sample $M^* \sim \mathbb{P}(\mathcal{M}^* \in \cdot \mid H_k)$ and passes it through the channel that achieves the rate-distortion limit (Equation 1) at this episode to get the M_k whose optimal policy is executed in the environment.

The core impetus for this work is to recognize that, for complex environments, pursuit of the exact MDP \mathcal{M}^* (as in PSRL) may be an entirely infeasible goal. Consider a MDP that represents control of a real-world, physical system; learning a transition function of the associated environment, at some level, demands that the agent internalize the laws of physics and motion with near-perfect accuracy. More formally, identifying \mathcal{M}^* demands the agent obtain exactly $\mathbb{H}_1(\mathcal{M}^*)$ bits of information from the environment which, under an uninformative prior, may either be prohibitively large by far exceeding the agent’s capacity constraints or be simply impractical under time and resource constraints.

As a remedy for this problem, we embrace the idea of being “sufficiently satisfying” or *satisficing* [148, 140, 138, 139, 15, 16]; as succinctly stated by Herbert A. Simon during his 1978 Nobel Memorial Lecture, “decision makers can satisfice either by finding optimum solutions for a simplified world, or by finding satisfactory solutions for a more realistic world.” Rather than spend an inordinate amount of time trying to recover an optimum solution to the true environment, we will instead design an algorithm that pursues optimum solutions for a sequence of simplified environments. In the next section, our analysis demonstrates that finding such optimum solutions for simplified worlds ultimately acts as a mechanism for achieving a satisfactory solution for the realistic, complex world. Naturally, the loss of fidelity between the simplified and true environments translates into a fixed amount of regret that an agent designer consciously and willingly accepts for two reasons: (1) they expect a reduction in the amount of time, data, and bits of information needed to identify the simplified environment and (2) in tasks where the environment encodes irrelevant information and

exact knowledge is not needed to achieve optimal behavior [66, 72, 73, 161], this worst-case error term may be negligible anyways while still maintaining greater efficiency than traditional PSRL.

Recalling Example 1 that revolves around a particular sequence of MDPs, $\{\mathcal{M}_n\}_{n \in [N]}$, we note that as the distortion threshold D increases, the significance of MDPs in the sequence indexed by larger values of $n \in [N]$ rapidly diminishes. As $D \uparrow \infty$, the lossy compression $\widetilde{\mathcal{M}}_k$ needn't convey information about any of the MDPs in $\{\mathcal{M}_n\}_{n \in [N]}$. Conversely, at $D = 0$, a VSRL agent must necessarily obtain enough information about the entire sequence so as to facilitate planning over Π and \mathcal{V} . In between, however, the agent need only concern itself with a particular subsequence of $\{\mathcal{M}_n\}_{n \in [N]}$ while the remaining MDPs can be ignored due to their negligible contribution to overall value and, therefore, expected distortion under $d_{\Pi, \mathcal{V}}$.

5 Regret Analysis

In this section, we offer an information-theoretic analysis of VSRL (Algorithm 2) before refining our regret bounds to the tabular setting. We conclude by highlighting how our performance guarantees can be expressed via a notion of agent capacity that is considerate of real-world reinforcement learning.

5.1 An Information-Theoretic Bayesian Regret Bound

To establish a Bayesian regret bound for VSRL we first require a regret decomposition that acknowledges the agent's new objective of identifying an approximately value-equivalent MDP in each episode, $\widetilde{\mathcal{M}}_k$, rather than the true MDP \mathcal{M}^* . Crucially, this regret decomposition leverages the precise form of our distortion function $d_{\Pi, \mathcal{V}}(\mathcal{M}^*, \widetilde{\mathcal{M}}_k)$.

Theorem 1. *Take any $\Pi \supseteq \{\mathcal{S} \rightarrow \mathcal{A}\}$, any $\mathcal{V} \supseteq \{V^\pi \mid \pi \in \Pi^H\}$, and fix any $D \geq 0$. For each episode $k \in [K]$, let $\widetilde{\mathcal{M}}_k$ be any MDP that achieves the rate-distortion limit of $\mathcal{R}_k^{\Pi, \mathcal{V}}(D)$ with information source $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_k)$ and distortion function $d_{\Pi, \mathcal{V}}$. Then,*

$$\text{BAYESREGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}) \leq \mathbb{E} \left[\sum_{k=1}^K \mathbb{E}_k \left[V_{\widetilde{\mathcal{M}}_k, 1}^* - V_{\widetilde{\mathcal{M}}_k, 1}^{\pi^{(k)}} \right] \right] + 2KH\sqrt{D}.$$

Theorem 1 shows how the Bayesian regret incurred by VSRL can be separated into an error term the agent must pay for learning a simplified MDP $\widetilde{\mathcal{M}}_k$, rather than \mathcal{M}^* , and the Bayesian regret incurred while trying to learn $\widetilde{\mathcal{M}}_k$. This first term mirrors the satisficing regret of Russo and Van Roy [138, 139] for multi-armed bandits where the performance of the agent in the k th episode is being measured with respect to a compressed MDP $\widetilde{\mathcal{M}}_k$, rather than the true MDP \mathcal{M}^* . While further discussion on the choices of Π and \mathcal{V} is provided later in this section, we simply note that the conditions placed upon them in Theorem 1 are an artifact of VSRL only executing optimal policies in each time period $h \in [H]$ which, under the assumptions of our problem formulation, are deterministic.

The remainder of this section is devoted to an analysis for establishing an information-theoretic bound on the satisficing regret term of Theorem 1. A central tool of our analysis will be the information ratio [136, 137] at the k th episode:

$$\Gamma_k \triangleq \frac{\mathbb{E}_k \left[V_{\widetilde{\mathcal{M}}_k, 1}^* - V_{\widetilde{\mathcal{M}}_k, 1}^{\pi^{(k)}} \right]^2}{\mathbb{I}_k(\widetilde{\mathcal{M}}_k; \tau_k, M_k)} \quad \forall k \in [K].$$

In words, the information ratio is the ratio between squared expected regret in the k th episode with respect to $\widetilde{\mathcal{M}}_k$ and the information gained about $\widetilde{\mathcal{M}}_k$ in the k th episode by sampling MDP M_k and observing trajectory τ_k , given the current history H_k . Numerous prior works have leveraged similar or generalized types of information ratios for analyzing multi-armed bandit problems [135–139, 54, 99, 169, 36, 15, 98] as well as reinforcement-learning problems [104]; in comparison to the latter, we simply note that our analysis bears stronger resemblance to those in multi-armed bandits by not constructing confidence sets over MDPs [121, 120, 104], avoiding a restricted focus to tabular problems. That said, our results are contingent upon the existence of a uniform upper bound to the information ratios across all episodes, a non-trivial result [78] that we leave to future work.

Through our information-ratio analysis, we obtain the following information-theoretic bound on satisficing Bayesian regret:

Theorem 2. If $\Gamma_k \leq \bar{\Gamma}$, for all $k \in [K]$, then $\mathbb{E} \left[\sum_{k=1}^K \mathbb{E}_k \left[V_{\widetilde{\mathcal{M}}_{k,1}}^* - V_{\widetilde{\mathcal{M}}_{k,1}}^{\pi^{(k)}} \right] \right] \leq \sqrt{\bar{\Gamma} K \mathcal{R}_1^{\Pi, \mathcal{V}}(D)}$.

An immediate consequence of the preceding theorems is the following corollary which establishes our main result, an information-theoretic Bayesian regret bound for VSRL. We omit the proof as it follows directly from applying Theorems 1 and 2 in sequence.

Corollary 1. Take any $\Pi \supseteq \{\mathcal{S} \rightarrow \mathcal{A}\}$, any $\mathcal{V} \supseteq \{V^\pi \mid \pi \in \Pi^H\}$, and fix any $D > 0$. For any prior distribution $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_1)$, if $\Gamma_k \leq \bar{\Gamma}$ for all $k \in [K]$, then VSRL (Algorithm 2) has $\text{BAYESREGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}) \leq \sqrt{\bar{\Gamma} K \mathcal{R}_1^{\Pi, \mathcal{V}}(D)} + 2KH\sqrt{D}$.

Once again we recall that, since the rate-distortion function is well-defined for arbitrary source and channel output random variables defined on abstract alphabets [43], the Bayesian regret bound of Corollary 1 holds for any finite-horizon, episodic MDP, extending beyond past analyses of PSRL constrained only to tabular MDPs. We defer a discussion of practical considerations for implementing VSRL to the appendix.

At this point, we call attention to the parameterization of our lossy compression problem by a particular policy class Π and value function class \mathcal{V} , whose dependence we inherit from the value equivalence principle [72]. The next result clarifies how the performance of VSRL is affected by fluctuations in these classes via a dominance relationship [150] between the induced distortion functions.

Lemma 1. For any two Π, Π' and any $\mathcal{V}, \mathcal{V}'$ such that $\Pi' \subseteq \Pi \subseteq \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}$ and $\mathcal{V}' \subseteq \mathcal{V} \subseteq \{\mathcal{S} \rightarrow \mathbb{R}\}$, we have $\mathcal{R}_k^{\Pi, \mathcal{V}}(D) \geq \mathcal{R}_k^{\Pi', \mathcal{V}'}(D)$, $\forall k \in [K], D > 0$.

Property 3 of Grimm et al. [72] highlights how the set of value-equivalent MDPs grows as the policy and value function classes shrink. Lemma 1 provides an intuitive, information-theoretic counterpart to their result where, as the sets of policies and value functions over which models will be assessed diminish, an agent may naturally compress more aggressively and throw away larger quantities of bits from each source distribution over the true MDP \mathcal{M}^* .

Since a compressed MDP $\widetilde{\mathcal{M}}_k$ that achieves the rate-distortion limit has *expected* distortion bounded by D , one may wonder how the probability of not recovering an approximately-value-equivalent MDP scales as $D \uparrow \infty$. To that end, we conclude this section with a final result that brings clarity to this via a generalization [60] of Fano's inequality [64]. We leave investigation of other generalizations of Fano's inequality that might yield similarly interesting results to future work [160, 8].

Lemma 2. Take any $\Pi \subseteq \{\mathcal{S} \rightarrow \Delta(\mathcal{A})\}$ and $\mathcal{V} \subseteq \{\mathcal{S} \rightarrow \mathbb{R}\}$. For any $D \geq 0$ and any $k \in [K]$, define $\delta = \sup_{\widetilde{\mathcal{M}} \in \mathfrak{M}} \mathbb{P}(d_{\Pi, \mathcal{V}}(\mathcal{M}^*, \widetilde{\mathcal{M}}) \leq D \mid H_k)$. Then,

$$\sup_{\widetilde{\mathcal{M}} \in \Lambda_k(D)} \mathbb{P}(d_{\Pi, \mathcal{V}}(\mathcal{M}^*, \widetilde{\mathcal{M}}) > D \mid H_k) \geq 1 - \frac{\mathcal{R}_k^{\Pi, \mathcal{V}}(D) + \log(2)}{\log\left(\frac{1}{\delta}\right)}.$$

For any episode $k \in [K]$, the left-hand side of the inequality in Lemma 2 denotes the worst-case error probability of sampling a compressed MDP $\widetilde{\mathcal{M}}$ that is not approximately-value-equivalent to \mathcal{M}^* . The right-hand side conveys that, in order to avoid such an error with reasonable probability, one requires a setting of $D < \infty$ such that $\mathcal{R}_k^{\Pi, \mathcal{V}}(D) \approx \log\left(\frac{1}{\delta}\right)$.

5.2 Specializing to Tabular MDPs

While the preceding subsection constitutes the main contribution of this paper, the presence of information-theoretic terms makes it difficult to compare our guarantees to those obtained in prior work, which typically focuses on the tabular setting. To help remedy this, we offer the following theorem which restricts focus to the case where the agent pursues an exactly value-equivalent model of the tabular environment. Notably, the results of this section still retain a dependence on a uniform upper bound to the information ratio whose exact form is a result left to future work.

Theorem 3. Take any $\Pi \supseteq \{\mathcal{S} \rightarrow \mathcal{A}\}$, any $\mathcal{V} \supseteq \{V^\pi \mid \pi \in \Pi^H\}$, and let $D = 0$. For any prior distribution $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_1)$ over tabular MDPs, if $\Gamma_k \leq \bar{\Gamma}$ for all $k \in [K]$, then VSRL (Algorithm 2) has $\text{BAYESREGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}) \leq \mathcal{O}\left(|\mathcal{S}| \sqrt{\bar{\Gamma}} |\mathcal{A}| K\right)$.

An immediate observation is that the Bayesian regret bound of Theorem 3 matches the dependence on the number of states, $|\mathcal{S}|$, obtained in the first (weaker) guarantee established for PSRL by Osband et al. [121]; we suspect that this guarantee for VSRL is unimprovable without further distributional assumptions [120, 119]. As an alternative, we contemplate how a change in the distortion measure used by VSRL might incur an improved regret bound when specialized to the tabular setting.

Specifically, notice that the only piece of the VSRL analysis tethered to the particular form of the distortion function $d_{\Pi, \nu}(\mathcal{M}, \widehat{\mathcal{M}})$ is Theorem 1, while all other components remain agnostic to the precise criterion for assessing the loss of fidelity between original and compressed MDPs. Consequently, there is potential for a modified distortion function to offer an improved regret analysis relative to Theorem 3. Rather than concerning ourselves with planning over multiple behaviors, we consider a distortion function based solely on the optimal action-value functions:

$$d_{Q^*}(\mathcal{M}, \widehat{\mathcal{M}}) = \sup_{h \in [H]} \|Q_{\mathcal{M}, h}^* - Q_{\widehat{\mathcal{M}}, h}^*\|_{\infty}^2 = \sup_{h \in [H]} \sup_{(s, a) \in \mathcal{S} \times \mathcal{A}} |Q_{\mathcal{M}, h}^*(s, a) - Q_{\widehat{\mathcal{M}}, h}^*(s, a)|^2.$$

We use $\mathcal{R}_k^{Q^*}(D)$ to denote the rate-distortion function under this new measure of distortion, $d_{Q^*}(\mathcal{M}, \widehat{\mathcal{M}})$. In order for this new distortion function to be compatible with VSRL, we require an analogue to the regret decomposition of Theorem 1.

Theorem 4. Fix any $D \geq 0$ and, for each episode $k \in [K]$, let $\widetilde{\mathcal{M}}_k$ be any MDP that achieves the rate-distortion limit of $\mathcal{R}_k^{Q^*}(D)$ with information source $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_k)$ and distortion function d_{Q^*} .

$$\text{Then, } \text{BAYESREGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}) \leq \mathbb{E} \left[\sum_{k=1}^K \mathbb{E}_k \left[V_{\widetilde{\mathcal{M}}_k, 1}^* - V_{\widetilde{\mathcal{M}}_k, 1}^{\pi^{(k)}} \right] \right] + 2K(H+1)\sqrt{D}.$$

With this regret decomposition in hand, we may recover the analogue to Corollary 1, whose proof is immediate and, therefore, omitted.

Corollary 2. Fix any $D > 0$. For any prior distribution $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_1)$, if $\Gamma_k \leq \bar{\Gamma}$ for all $k \in [K]$, then VSRL (Algorithm 2) with distortion function d_{Q^*} has $\text{BAYESREGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}) \leq \sqrt{\bar{\Gamma} K \mathcal{R}_1^{Q^*}(D)} + 2K(H+1)\sqrt{D}$.

As illustrated by the following lemma, the significance of this change in distortion measure from $d_{\Pi, \nu}$ to d_{Q^*} is that the optimal action-value functions may now act as an information bottleneck [158] between the original MDP \mathcal{M}^* and compressed MDP $\widetilde{\mathcal{M}}_k$.

Lemma 3. For each episode $k \in [K]$ and for $D = 0$, let $\widetilde{\mathcal{M}}_k$ be any MDP that achieves the rate-distortion limit of $\mathcal{R}_k^{Q^*}(D)$ with information source $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_k)$ and distortion function d_{Q^*} . Then, we have the Markov chain $\mathcal{M}^* \rightarrow Q_{\mathcal{M}^*}^* \rightarrow \widetilde{\mathcal{M}}_k$, where $Q_{\mathcal{M}^*}^* = \{Q_{\mathcal{M}^*, h}^*\}_{h \in [H]}$ is the collection of random variables denoting the optimal action-value functions of \mathcal{M}^* .

Lemma 3, through the data-processing inequality, immediately leads us to an analogue of Theorem 3 that matches the dependence on $|\mathcal{S}|$ in the best known Bayesian regret bound for PSRL [120].

Theorem 5. For $D = 0$ and any prior distribution $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_1)$ over tabular MDPs, if $\Gamma_k \leq \bar{\Gamma}$ for all $k \in [K]$, then VSRL with distortion function d_{Q^*} has $\text{BAYESREGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}) \leq \tilde{\mathcal{O}} \left(\sqrt{\bar{\Gamma} |\mathcal{S}| |\mathcal{A}| KH} \right)$.

Ultimately, Theorem 5 confirms that while there is great flexibility in the original definition of value equivalence to support planning across multiple policies and value functions, focusing on optimal value functions gives rise to more efficient learning. Moreover, comparing the result with the PSRL regret bound of Osband and Van Roy [120] for tabular MDPs, this suggests an achievable uniform upper bound to the information ratio as $\bar{\Gamma} \lesssim H^2$, where the \lesssim accounts for numerical constants and logarithmic factors.

5.3 Capacity-Sensitive Performance Guarantees

We recognize that the information-theoretic regret bounds of the previous two sections, like many other guarantees for provably-efficient reinforcement learning before them, implicitly and unrealistically

assume that an agent is of unbounded capacity and may pursue any approximately-value-equivalent model under a given distortion threshold D . In the context of real-world reinforcement learning [62, 105], however, fundamental limits on computational resources and time leave an agent designer with a bounded agent to be deployed within an overwhelmingly complex environment. As such, this designer may seldom be in a position to dictate an ideal or desired sub-optimality threshold D , but rather must make do with a known constraint on agent capacity; guarantees on sample-efficient reinforcement learning cognizant of such a fundamental constraint are nascent.

While there are numerous possibilities for how one might choose to formally characterize agent capacity, we here adopt a fundamental perspective that learning is the process of acquiring information and so take this capacity to imply the existence of a non-negative real value $R \in \mathbb{R}_{>0}$ such that the agent may only acquire and retain exactly R bits of information. To help contextualize this notion of agent capacity, we introduce the distortion-rate function [146, 25, 41] which quantifies the fundamental limit of expected distortion under an information constraint:

$$\mathcal{D}_k^{Q^*}(R) = \inf_{\widetilde{\mathcal{M}} \in \Upsilon_k(R)} \mathbb{E}_k \left[d_{Q^*}(\mathcal{M}^*, \widetilde{\mathcal{M}}) \right] \quad \mathcal{D}_k^{Q^*}(R) = \inf_{\widetilde{\mathcal{M}} \in \Upsilon_k(R)} \mathbb{E}_k \left[d_{Q^*}(\mathcal{M}^*, \widetilde{\mathcal{M}}) \right], \quad (2)$$

where the infimum is taken over all channels with bounded rate, $\Upsilon_k(R) \triangleq \{\widetilde{\mathcal{M}} : \Omega \rightarrow \mathfrak{M} \mid \mathbb{I}_k(\mathcal{M}^*; \widetilde{\mathcal{M}}) \leq R\}$. In words, given the agent's current beliefs over the true MDP $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_k)$, the infimum of the distortion-rate function is taken over all potential lossy compressions of the environment that fall within the agent's capacity constraint of R bits and identifies the one that preserves the most useful information, as measured by the distortion function. Conveniently, the rate-distortion function and distortion-rate function are inverses of one another [41] ($\mathcal{R}(\mathcal{D}(R)) = R$) such that we recover the following two capacity-sensitive regret bounds directly from Corollaries 1 and 2 by simply taking the input distortion threshold of VSRL equal to the associated distortion-rate function in the first episode ($D = \mathcal{D}_1^{\Pi, \mathcal{V}}(R)$ and $D = \mathcal{D}_1^{Q^*}(R)$, respectively).

Corollary 3. *Take any $\Pi \supseteq \{\mathcal{S} \rightarrow \mathcal{A}\}$, any $\mathcal{V} \supseteq \{V^\pi \mid \pi \in \Pi^H\}$, and let $R > 0$ be the agent capacity. For any prior distribution $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_1)$, if $\Gamma_k \leq \bar{\Gamma}$ for all $k \in [K]$, then VSRL (Algorithm 2) with distortion function $d_{\Pi, \mathcal{V}}$ has $\text{BAYESREGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}) \leq \sqrt{\bar{\Gamma}KR} + 2KH\sqrt{\mathcal{D}_1^{\Pi, \mathcal{V}}(R)}$.*

Corollary 4. *Let $R > 0$ be the agent capacity. For any prior distribution $\mathbb{P}(\mathcal{M}^* \in \cdot \mid H_1)$, if $\Gamma_k \leq \bar{\Gamma}$ for all $k \in [K]$, then VSRL (Algorithm 2) with distortion function d_{Q^*} has $\text{BAYESREGRET}(K, \pi^{(1)}, \dots, \pi^{(K)}) \leq \sqrt{\bar{\Gamma}KR} + 2K(H+1)\sqrt{\mathcal{D}_1^{Q^*}(R)}$.*

Turning back to Example 1, note how an agent with significantly limited capacity cannot possibly hope to capture all the granularity contained in the entire MDP sequence $\{\mathcal{M}_n\}_{n \in [N]}$, for large values of N . For a capacity of exactly R bits, Corollaries 3 and 4 immediately translate this fundamental limit into a corresponding performance guarantee, allowing the agent to identify a subsequence $\{\mathcal{M}_n\}_{n \in [C]}$ for some $C \ll N$ which only requires gathering R bits of information from the environment.

6 Conclusion

In this paper, we began with a finite-horizon, episodic MDP and considered the ramifications of a real-world reinforcement-learning scenario wherein the relative complexity of the environment is so immense that an agent may find itself incapable of perfectly recovering optimal behavior. An immediate consequence of this reality is the need to strike an appropriate balance between what is performant and what is achievable. We introduced the VSRL algorithm for incrementally synthesizing *simple* and *useful* approximations of the environment from which an agent might still recover near-optimal behaviors. Recognizing the information-theoretic nature of this lossy MDP compression, we provided an analysis of VSRL whose performance guarantees, by virtue of rate-distortion theory, are twofold. The first set of guarantees ensure VSRL recovers the simplest compression of the environment which still incurs bounded sub-optimality, as specified by the agent designer. Alternatively, the second set of guarantees maintain that VSRL finds the best compression of the environment subject to constraints on agent capacity. Through our general problem formulation and information-theoretic analysis, both regret bounds hold for any finite-horizon, episodic MDP, regardless of whether or not the state-action space is finite. That said, the question of how to practically instantiate VSRL for high-dimensional settings is an open problem left to future work.

Acknowledgements

The authors gratefully acknowledge Christopher Grimm for initial discussions that provided an impetus for this work. The authors also thank Adithya Devraj, Shi Dong, John Duchi, Hong Jun Jeon, Saurabh Kumar, and Xiuyuan (Lucy) Lu for insightful comments on various components of the paper. Financial support from Army Research Office (ARO) grant W911NF2010055 is gratefully acknowledged.

References

- [1] Romina Abachi, Mohammad Ghavamzadeh, and Amir-massoud Farahmand. Policy-aware model learning for policy gradient methods. *arXiv preprint arXiv:2003.00030*, 2020. 5
- [2] Yasin Abbasi-Yadkori and Csaba Szepesvari. Bayesian optimal control of smoothly parameterized systems: The lazy posterior sampling algorithm. *arXiv preprint arXiv:1406.3926*, 2014. 22
- [3] David Abel. *A Theory of Abstraction in Reinforcement Learning*. PhD thesis, Brown University, 2020. 23
- [4] David Abel, David Hershkowitz, and Michael Littman. Near optimal behavior via approximate state abstraction. In *International Conference on Machine Learning*, pages 2915–2923. PMLR, 2016. 23, 25
- [5] David Abel, Dilip Arumugam, Lucas Lehnert, and Michael Littman. State abstractions for lifelong reinforcement learning. In *International Conference on Machine Learning*, pages 10–19, 2018.
- [6] David Abel, Dilip Arumugam, Kavosh Asadi, Yuu Jinnai, Michael L Littman, and Lawson LS Wong. State abstraction as compression in apprenticeship learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3134–3142, 2019. 22, 23, 25
- [7] David Abel, Nate Umbanhowar, Khimya Khetarpal, Dilip Arumugam, Doina Precup, and Michael Littman. Value preserving state-action abstractions. In *International Conference on Artificial Intelligence and Statistics*, pages 1639–1650. PMLR, 2020. 23
- [8] Shuchin Aeron, Venkatesh Saligrama, and Manqi Zhao. Information theoretic bounds for compressed sensing. *IEEE Transactions on Information Theory*, 56(10):5111–5130, 2010. 8
- [9] Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. FLAMBE: Structural complexity and representation learning of low rank MDPs. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20095–20107. Curran Associates, Inc., 2020. 23
- [10] Shipra Agrawal and Randy Jia. Optimistic posterior sampling for reinforcement learning: worst-case regret bounds. In *Advances in Neural Information Processing Systems*, pages 1184–1194, 2017. 1, 22
- [11] Alexander Alemi, Ben Poole, Ian Fischer, Joshua Dillon, Rif A Saurous, and Kevin Murphy. Fixing a broken ELBO. In *International Conference on Machine Learning*, pages 159–168. PMLR, 2018. 25
- [12] Mauricio Araya-López, Vincent Thomas, and Olivier Buffet. Near-optimal BRL using optimistic local transitions. In *Proceedings of the 29th International Conference on Machine Learning*, pages 515–522, 2012. 22
- [13] Suguru Arimoto. An algorithm for computing the capacity of arbitrary discrete memoryless channels. *IEEE Transactions on Information Theory*, 18(1):14–20, 1972. 25
- [14] Dilip Arumugam and Benjamin Van Roy. Randomized value functions via posterior state-abstraction sampling. *arXiv preprint arXiv:2010.02383*, 2020. 23

- [15] Dilip Arumugam and Benjamin Van Roy. Deciding what to learn: A rate-distortion approach. In *International Conference on Machine Learning*, pages 373–382. PMLR, 2021. 2, 5, 6, 7, 22, 25, 29
- [16] Dilip Arumugam and Benjamin Van Roy. The value of information when deciding what to learn. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 6, 22, 25
- [17] Kavosh Asadi, Dipendra Misra, and Michael Littman. Lipschitz continuity in model-based reinforcement learning. In *International Conference on Machine Learning*, pages 264–273. PMLR, 2018. 5
- [18] 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, 2009. 22
- [19] Karl Johan Åström. Optimal control of Markov processes with incomplete state information. *Journal of Mathematical Analysis and Applications*, 10(1):174–205, 1965. 23
- [20] Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 89–96, 2009. 1, 22
- [21] Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pages 463–474. PMLR, 2020. 5
- [22] Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pages 263–272. PMLR, 2017. 1, 22
- [23] Peter L Bartlett and Ambuj Tewari. REGAL: a regularization based algorithm for reinforcement learning in weakly communicating MDPs. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 35–42, 2009. 1, 22
- [24] Richard Bellman. A Markovian decision process. *Journal of Mathematics and Mechanics*, pages 679–684, 1957. 3
- [25] Toby Berger. *Rate Distortion Theory: A Mathematical Basis for Data Compression*. Prentice-Hall, 1971. 2, 4, 10
- [26] Toby Berger and Jerry D Gibson. Lossy source coding. *IEEE Transactions on Information Theory*, 44(6):2693–2723, 1998. 4
- [27] Donald A. Berry, Robert W. Chen, Alan Zame, David C. Heath, and Larry A. Shepp. Bandit problems with infinitely many arms. *Ann. Statist.*, 25(5):2103–2116, 10 1997. 22
- [28] Dimitri Bertsekas and Steven E Shreve. *Stochastic optimal control: the discrete-time case*, volume 5. Athena Scientific, 1996. 4
- [29] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 1995. 5
- [30] Dimitri P. Bertsekas and David A. Castañón. Adaptive aggregation methods for infinite horizon dynamic programming. *IEEE Transactions on Automatic Control*, 34(6):589–598, 1989. 23
- [31] Richard Blahut. Computation of channel capacity and rate-distortion functions. *IEEE Transactions on Information Theory*, 18(4):460–473, 1972. 25
- [32] Thomas Bonald and Alexandre Proutiere. Two-target algorithms for infinite-armed bandits with Bernoulli rewards. In *Advances in Neural Information Processing Systems*, pages 2184–2192, 2013. 22
- [33] Vivek S Borkar, Sanjoy Mitter, and Sekhar Tatikonda. Markov control problems under communication constraints. *Communications in Information and Systems*, 1(1):15–32, 2001. 22

- [34] Pinhas Boukris. An upper bound on the speed of convergence of the Blahut algorithm for computing rate-distortion functions (corresp.). *IEEE Transactions on Information Theory*, 19(5):708–709, 1973. [25](#)
- [35] Ronen I Brafman and Moshe Tennenholtz. R-MAX—a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3(Oct):213–231, 2002. [1](#), [22](#)
- [36] Sébastien Bubeck and Mark Sellke. First-order Bayesian regret analysis of Thompson sampling. In *Algorithmic Learning Theory*, pages 196–233. PMLR, 2020. [7](#)
- [37] Sébastien Bubeck, Rémi Munos, Gilles Stoltz, and Csaba Szepesvári. X-armed bandits. *Journal of Machine Learning Research*, 12(5), 2011. [22](#)
- [38] Sébastien Bubeck, Nicolò Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012. [2](#), [22](#)
- [39] Pablo Samuel Castro and Doina Precup. Using linear programming for Bayesian exploration in Markov decision processes. In *IJCAI*, volume 24372442, 2007. [22](#)
- [40] Mung Chiang and Stephen Boyd. Geometric programming duals of channel capacity and rate distortion. *IEEE Transactions on Information Theory*, 50(2):245–258, 2004. [25](#)
- [41] Thomas M Cover and Joy A Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012. [3](#), [4](#), [10](#), [30](#), [33](#), [35](#)
- [42] Imre Csiszár. On the computation of rate-distortion functions (corresp.). *IEEE Transactions on Information Theory*, 20(1):122–124, 1974. [25](#)
- [43] Imre Csiszár. On an extremum problem of information theory. *Studia Scientiarum Mathematicarum Hungarica*, 9, 1974. [5](#), [8](#), [25](#)
- [44] Brandon Cui, Yinlam Chow, and Mohammad Ghavamzadeh. Control-aware representations for model-based reinforcement learning. In *International Conference on Learning Representations*, 2020. [5](#)
- [45] Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 2*, pages 2818–2826, 2015. [1](#), [22](#)
- [46] Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying PAC and regret: uniform PAC bounds for episodic reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5717–5727, 2017. [1](#), [22](#)
- [47] Christoph Dann, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. On oracle-efficient PAC RL with rich observations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 1429–1439, 2018. [22](#)
- [48] Justin Dauwels. Numerical computation of the capacity of continuous memoryless channels. In *Proceedings of the 26th Symposium on Information Theory in the BENELUX*, pages 221–228. Citeseer, 2005. [25](#)
- [49] Peter Dayan. Improving generalization for temporal difference learning: The successor representation. *Neural computation*, 5(4):613–624, 1993. [26](#)
- [50] Thomas Dean and Robert Givan. Model minimization in Markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 106–111. AAAI Press, 1997. [23](#)
- [51] Richard Dearden, Nir Friedman, and Stuart Russell. Bayesian Q-learning. In *Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence*, pages 761–768, 1998. [22](#)

- [52] Richard Dearden, Nir Friedman, and David Andre. Model based Bayesian exploration. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, pages 150–159, 1999. 22
- [53] Yash Deshpande and Andrea Montanari. Linear bandits in high dimension and recommendation systems. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1750–1754. IEEE, 2012. 22
- [54] Shi Dong and Benjamin Van Roy. An information-theoretic analysis for Thompson sampling with many actions. In *Advances in Neural Information Processing Systems*, pages 4157–4165, 2018. 7, 22
- [55] Shi Dong, Benjamin Van Roy, and Zhengyuan Zhou. Provably efficient reinforcement learning with aggregated states. *arXiv preprint arXiv:1912.06366*, 2019. 23
- [56] Shi Dong, Benjamin Van Roy, and Zhengyuan Zhou. Simple agent, complex environment: Efficient reinforcement learning with agent state. *arXiv preprint arXiv:2102.05261*, 2021. 1, 22, 23
- [57] Pierluca D’Oro, Alberto Maria Metelli, Andrea Tirinzoni, Matteo Papini, and Marcello Restelli. Gradient-aware model-based policy search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3801–3808, 2020. 5
- [58] Simon Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudik, and John Langford. Provably efficient RL with rich observations via latent state decoding. In *International Conference on Machine Learning*, pages 1665–1674. PMLR, 2019. 22, 23
- [59] John C. Duchi. *Lecture Notes for Statistics 311/Electrical Engineering 377, Stanford University*. 2021. 3, 33, 36
- [60] John C. Duchi and Martin J. Wainwright. Distance-based and continuum Fano inequalities with applications to statistical estimation. *arXiv preprint arXiv:1311.2669*, 2013. 8, 33
- [61] Michael O’Gordon Duff. *Optimal Learning: Computational procedures for Bayes-adaptive Markov decision processes*. University of Massachusetts Amherst, 2002. 22
- [62] Gabriel Dulac-Arnold, Nir Levine, Daniel J Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester. Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Machine Learning*, pages 1–50, 2021. 10
- [63] Vikranth Dwaracherla, Xiuyuan Lu, Morteza Ibrahimi, Ian Osband, Zheng Wen, and Benjamin Van Roy. Hypermodels for exploration. In *International Conference on Learning Representations*, 2020. 2, 25
- [64] Robert M. Fano. *Class Notes for MIT Course 6.574: Transmission of Information, MIT, Cambridge, MA*. 1952. 8, 33
- [65] Amir-massoud Farahmand. Iterative value-aware model learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 9090–9101, 2018. 5
- [66] Amir-massoud Farahmand, Andre Barreto, and Daniel Nikovski. Value-aware loss function for model-based reinforcement learning. In *Artificial Intelligence and Statistics*, pages 1486–1494. PMLR, 2017. 5, 7
- [67] Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite Markov Decision Processes. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pages 162–169, 2004. 23
- [68] Norman Ferns, Pablo Samuel Castro, Doina Precup, and Prakash Panangaden. Methods for computing state similarity in Markov Decision Processes. *arXiv preprint arXiv:1206.6836*, 2012. 23

- [69] Sebastian Flennerhag, Jane X Wang, Pablo Sprechmann, Francesco Visin, Alexandre Galashov, Steven Kapturowski, Diana L Borsa, Nicolas Heess, Andre Barreto, and Razvan Pascanu. Temporal difference uncertainties as a signal for exploration. *arXiv preprint arXiv:2010.02255*, 2020. [26](#)
- [70] Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, and Aviv Tamar. Bayesian reinforcement learning: A survey. *Foundations and Trends® in Machine Learning*, 8(5-6):359–483, 2015. [3](#), [22](#)
- [71] Robert M. Gray. *Entropy and information theory*. Springer Science & Business Media, 2011. [3](#)
- [72] Christopher Grimm, Andre Barreto, Satinder Singh, and David Silver. The value equivalence principle for model-based reinforcement learning. *Advances in Neural Information Processing Systems*, 33, 2020. [5](#), [6](#), [7](#), [8](#)
- [73] Christopher Grimm, André Barreto, Greg Farquhar, David Silver, and Satinder Singh. Proper value equivalence. *Advances in Neural Information Processing Systems*, 34, 2021. [5](#), [7](#), [23](#)
- [74] Arthur Guez, David Silver, and Peter Dayan. Efficient Bayes-adaptive reinforcement learning using sample-based search. In *Proceedings of the 25th International Conference on Neural Information Processing Systems-Volume 1*, pages 1025–1033, 2012. [22](#)
- [75] Arthur Guez, David Silver, and Peter Dayan. Scalable and efficient Bayes-adaptive reinforcement learning based on Monte-Carlo tree search. *Journal of Artificial Intelligence Research*, 48:841–883, 2013.
- [76] Arthur Guez, Nicolas Heess, David Silver, and Peter Dayan. Bayes-adaptive simulation-based search with value function approximation. In *Advances in Neural Information Processing Systems*, pages 451–459, 2014. [22](#)
- [77] Assaf Hallak, Dotan Di Castro, and Shie Mannor. Contextual Markov decision processes. *arXiv preprint arXiv:1502.02259*, 2015. [25](#)
- [78] Botao Hao and Tor Lattimore. Regret bounds for information-directed reinforcement learning. *arXiv preprint arXiv:2206.04640*, 2022. [7](#)
- [79] Matthew T Harrison and Ioannis Kontoyiannis. Estimation of the rate–distortion function. *IEEE Transactions on Information Theory*, 54(8):3757–3762, 2008. [25](#)
- [80] Cassius T. Ionescu-Tulcea. Mesures dans les espaces produits. *Atti Acad. Naz. Lincei Rend. Cl Sci. Fis. Mat. Nat.*, 8(7), 1949. [4](#)
- [81] Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010. [1](#), [22](#)
- [82] David Janz, Jiri Hron, Przemysław Mazur, Katja Hofmann, José Miguel Hernández-Lobato, and Sebastian Tschiatschek. Successor uncertainties: exploration and uncertainty in temporal difference learning. *Advances in Neural Information Processing Systems*, 32, 2019. [26](#)
- [83] Nan Jiang, Alex Kulesza, and Satinder Singh. Abstraction selection in model-based reinforcement learning. In *International Conference on Machine Learning*, pages 179–188, 2015. [23](#)
- [84] Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low Bellman rank are PAC-learnable. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017. [22](#)
- [85] Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is Q -learning provably efficient? In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 4868–4878, 2018. [1](#), [22](#)
- [86] Nicholas K Jong and Peter Stone. State abstraction discovery from irrelevant state variables. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, pages 752–757, 2005. [23](#)

- [87] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996. [1](#)
- [88] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998. [23](#)
- [89] Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *Proceedings of the Nineteenth International Conference on Machine Learning*, pages 267–274, 2002. [37](#)
- [90] Sham Machandranath Kakade. *On the Sample Complexity of Reinforcement Learning*. PhD thesis, Gatsby Computational Neuroscience Unit, University College London, 2003. [1](#), [22](#)
- [91] Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine Learning*, 49(2-3):209–232, 2002. [1](#), [22](#)
- [92] Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Multi-armed bandits in metric spaces. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 681–690, 2008. [22](#)
- [93] Andrei Nikolaevich Kolmogorov and Vladimir Mikhailovich Tikhomirov. ε -entropy and ε -capacity of sets in function spaces. *Uspekhi Matematicheskikh Nauk*, 14(2):3–86, 1959. [35](#)
- [94] J Zico Kolter and Andrew Y Ng. Near-Bayesian exploration in polynomial time. In *Proceedings of the 26th Annual International Conference on machine Learning*, pages 513–520, 2009. [22](#)
- [95] Victoria Kostina and Babak Hassibi. Rate-cost tradeoffs in control. *IEEE Transactions on Automatic Control*, 64(11):4525–4540, 2019. [22](#)
- [96] Akshay Krishnamurthy, Alekh Agarwal, and John Langford. PAC reinforcement learning with rich observations. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 1848–1856, 2016. [22](#)
- [97] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985. [2](#)
- [98] Tor Lattimore and Andras Gyorgy. Mirror descent and the information ratio. In *Conference on Learning Theory*, pages 2965–2992. PMLR, 2021. [7](#)
- [99] Tor Lattimore and Csaba Szepesvári. An information-theoretic approach to minimax regret in partial monitoring. In *Conference on Learning Theory*, pages 2111–2139. PMLR, 2019. [7](#)
- [100] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020. [2](#), [4](#), [22](#)
- [101] Lihong Li, Thomas J Walsh, and Michael L Littman. Towards a unified theory of state abstraction for MDPs. *ISAIM*, 4:5, 2006. [23](#), [25](#)
- [102] Evan Z Liu, Aditi Raghunathan, Percy Liang, and Chelsea Finn. Decoupling exploration and exploitation for meta-reinforcement learning without sacrifices. In *International conference on machine learning*, pages 6925–6935. PMLR, 2021. [25](#)
- [103] Xiuyuan Lu and Benjamin Van Roy. Ensemble sampling. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 3260–3268, 2017. [25](#)
- [104] Xiuyuan Lu and Benjamin Van Roy. Information-theoretic confidence bounds for reinforcement learning. *Advances in Neural Information Processing Systems*, 32:2461–2470, 2019. [7](#)
- [105] Xiuyuan Lu, Benjamin Van Roy, Vikranth Dwaracherla, Morteza Ibrahimi, Ian Osband, and Zheng Wen. Reinforcement Learning, Bit by Bit. *arXiv preprint arXiv:2103.04047*, 2021. [1](#), [2](#), [10](#), [22](#), [23](#)

- [106] Gerald Matz and Pierre Duhamel. Information geometric formulation and interpretation of accelerated Blahut-Arimoto-type algorithms. In *Information theory workshop*, pages 66–70. IEEE, 2004. 25
- [107] Dipendra Misra, Mikael Henaff, Akshay Krishnamurthy, and John Langford. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. In *International Conference on Machine Learning*, pages 6961–6971. PMLR, 2020. 23
- [108] Sanjoy Mitter and Anant Sahai. Information and control: Witsenhausen revisited. In *Learning, Control and Hybrid Systems*, pages 281–293. Springer, 1999. 22
- [109] Sanjoy K Mitter. Control with limited information. *European Journal of Control*, 7(2-3): 122–131, 2001. 22
- [110] Suraj Nair, Silvio Savarese, and Chelsea Finn. Goal-aware prediction: Learning to model what matters. In *International Conference on Machine Learning*, pages 7207–7219. PMLR, 2020. 5
- [111] Ziad Naja, Florence Alberge, and Pierre Duhamel. Geometrical interpretation and improvements of the Blahut-Arimoto’s algorithm. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2505–2508. IEEE, 2009. 25
- [112] Urs Niesen, Devavrat Shah, and Gregory Wornell. Adaptive alternating minimization algorithms. In *2007 IEEE International Symposium on Information Theory*, pages 1641–1645. IEEE, 2007. 25
- [113] Evgenii Nikishin, Romina Abachi, Rishabh Agarwal, and Pierre-Luc Bacon. Control-oriented model-based reinforcement learning with implicit differentiation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022. 5
- [114] Brendan O’Donoghue, Ian Osband, Remi Munos, and Volodymyr Mnih. The uncertainty Bellman equation and exploration. In *International Conference on Machine Learning*, pages 3836–3845, 2018. 2, 26
- [115] Brendan O’Donoghue, Ian Osband, and Catalin Ionescu. Making sense of reinforcement learning and probabilistic inference. In *International Conference on Learning Representations*, 2020. 25
- [116] Junhyuk Oh, Satinder Singh, and Honglak Lee. Value prediction network. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6120–6130, 2017. 5, 23
- [117] Ian Osband. *Deep Exploration via Randomized Value Functions*. PhD thesis, Stanford University, 2016. 26
- [118] Ian Osband and Benjamin Van Roy. Model-based reinforcement learning and the Eluder dimension. *Advances in Neural Information Processing Systems*, 27, 2014. 22
- [119] Ian Osband and Benjamin Van Roy. Gaussian-Dirichlet posterior dominance in sequential learning. *arXiv preprint arXiv:1702.04126*, 2017. 9
- [120] Ian Osband and Benjamin Van Roy. Why is posterior sampling better than optimism for reinforcement learning? In *International Conference on Machine Learning*, pages 2701–2710. PMLR, 2017. 1, 2, 6, 7, 9, 22, 38
- [121] Ian Osband, Daniel Russo, and Benjamin Van Roy. (More) efficient reinforcement learning via posterior sampling. *Advances in Neural Information Processing Systems*, 26:3003–3011, 2013. 1, 2, 7, 9, 22, 26, 38
- [122] Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via Bootstrapped DQN. In *Advances in Neural Information Processing Systems*, pages 4026–4034, 2016. 25, 26
- [123] Ian Osband, Benjamin Van Roy, and Zheng Wen. Generalization and exploration via randomized value functions. In *International Conference on Machine Learning*, pages 2377–2386, 2016. 2

- [124] Ian Osband, John Aslanides, and Albin Cassirer. Randomized prior functions for deep reinforcement learning. *Advances in Neural Information Processing Systems*, 31, 2018. 25, 26
- [125] Ian Osband, Benjamin Van Roy, Daniel J Russo, and Zheng Wen. Deep exploration via randomized value functions. *Journal of Machine Learning Research*, 20(124):1–62, 2019. 2, 26
- [126] Ian Osband, Zheng Wen, Mohammad Asghari, Morteza Ibrahimi, Xiyuan Lu, and Benjamin Van Roy. Epistemic neural networks. *arXiv preprint arXiv:2107.08924*, 2021. 25, 26
- [127] Ian Osband, Zheng Wen, Seyed Mohammad Asghari, Vikranth Dwaracherla, Botao Hao, Morteza Ibrahimi, Dieterich Lawson, Xiyuan Lu, Brendan O’Donoghue, and Benjamin Van Roy. Evaluating predictive distributions: Does Bayesian deep learning work? *arXiv preprint arXiv:2110.04629*, 2021. 25
- [128] Hari Palaiyanur and Anant Sahai. On the uniform continuity of the rate-distortion function. In *2008 IEEE International Symposium on Information Theory*, pages 857–861. IEEE, 2008. 25
- [129] Yury Polyanskiy and Yihong Wu. Lecture notes on information theory. 2019. 3
- [130] Martin L. Puterman. *Markov Decision Processes—Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, 1994. 3
- [131] Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *International conference on machine learning*, pages 5331–5340. PMLR, 2019. 25
- [132] Kenneth Rose. A mapping approach to rate-distortion computation and analysis. *IEEE Transactions on Information Theory*, 40(6):1939–1952, 1994. 25
- [133] Jonathan Rubin, Ohad Shamir, and Naftali Tishby. Trading value and information in MDPs. In *Decision Making with Imperfect Decision Makers*, pages 57–74. Springer, 2012. 22
- [134] Paat Rusmevichientong and John N Tsitsiklis. Linearly parameterized bandits. *Mathematics of Operations Research*, 35(2):395–411, 2010. 22
- [135] Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. *Advances in Neural Information Processing Systems*, 27:1583–1591, 2014. 7, 22, 38
- [136] Daniel Russo and Benjamin Van Roy. An information-theoretic analysis of Thompson sampling. *The Journal of Machine Learning Research*, 17(1):2442–2471, 2016. 7
- [137] Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. *Operations Research*, 66(1):230–252, 2018. 7, 22
- [138] Daniel Russo and Benjamin Van Roy. Satisficing in time-sensitive bandit learning. *arXiv preprint arXiv:1803.02855*, 2018. 2, 5, 6, 7, 22
- [139] Daniel Russo and Benjamin Van Roy. Satisficing in time-sensitive bandit learning. *Mathematics of Operations Research*, 2022. 2, 5, 6, 7, 22
- [140] Daniel Russo, David Tse, and Benjamin Van Roy. Time-sensitive bandit learning and satisficing Thompson sampling. *arXiv preprint arXiv:1704.09028*, 2017. 2, 6, 22
- [141] Daniel J Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on Thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1):1–96, 2018. 1, 22
- [142] Ilya O Ryzhov, Warren B Powell, and Peter I Frazier. The knowledge gradient algorithm for a general class of online learning problems. *Operations Research*, 60(1):180–195, 2012. 22
- [143] Jossy Sayir. Iterating the Arimoto-Blahut algorithm for faster convergence. In *2000 IEEE International Symposium on Information Theory (Cat. No. 00CH37060)*, page 235. IEEE, 2000. 25

- [144] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering Atari, Go, Chess and Shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020. 5, 23, 25
- [145] Ehsan Shafieepoorfard, Maxim Raginsky, and Sean P Meyn. Rationally inattentive control of Markov processes. *SIAM Journal on Control and Optimization*, 54(2):987–1016, 2016. 22, 25
- [146] Claude E. Shannon. Coding theorems for a discrete source with a fidelity criterion. *IRE Nat. Conv. Rec., March 1959*, 4:142–163, 1959. 2, 4, 10, 22
- [147] David Silver, Hado Hasselt, Matteo Hessel, Tom Schaul, Arthur Guez, Tim Harley, Gabriel Dulac-Arnold, David Reichert, Neil Rabinowitz, Andre Barreto, et al. The Predictron: End-to-end learning and planning. In *International Conference on Machine Learning*, pages 3191–3199. PMLR, 2017. 5, 23
- [148] Herbert A. Simon. Models of bounded rationality. *Economic Analysis and Public Policy*, MIT Press, Cambridge, Mass, 1982. 6
- [149] Jonathan Sorg, Satinder Singh, and Richard L Lewis. Variance-based rewards for approximate Bayesian reinforcement learning. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, pages 564–571, 2010. 22
- [150] Jan-Erik Stjernvall. Dominance—a relation between distortion measures. *IEEE Transactions on Information Theory*, 29(6):798–807, 1983. 8, 32, 33
- [151] Alexander L Strehl, Lihong Li, and Michael L Littman. Reinforcement learning in finite MDPs: PAC analysis. *Journal of Machine Learning Research*, 10(Nov):2413–2444, 2009. 1, 22
- [152] Malcolm JA Strens. A Bayesian framework for reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pages 943–950, 2000. 2, 6, 22
- [153] Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Model-based RL in contextual decision processes: PAC bounds and exponential improvements over model-free approaches. In *Conference on Learning Theory*, pages 2898–2933. PMLR, 2019. 22
- [154] Richard S Sutton and Andrew G Barto. Introduction to reinforcement learning. 1998. 1
- [155] Sekhar Tatikonda and Sanjoy Mitter. Control under communication constraints. *IEEE Transactions on Automatic Control*, 49(7):1056–1068, 2004. 22
- [156] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933. 1, 22
- [157] Naftali Tishby and Daniel Polani. Information theory of decisions and actions. In *Perception-action cycle*, pages 601–636. Springer, 2011. 22
- [158] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000. 9, 25
- [159] Benjamin Van Roy. Performance loss bounds for approximate value iteration with state aggregation. *Mathematics of Operations Research*, 31(2):234–244, 2006. 23, 25
- [160] Sergio Verdú et al. Generalizing the Fano inequality. *IEEE Transactions on Information Theory*, 40(4):1247–1251, 1994. 8
- [161] Claas A Voelcker, Victor Liao, Animesh Garg, and Amir-massoud Farahmand. Value gradient weighted model-based reinforcement learning. In *International Conference on Learning Representations*, 2022. 5, 7
- [162] Pascal O Vontobel, Aleksandar Kavcic, Dieter M Arnold, and Hans-Andrea Loeliger. A generalization of the Blahut–Arimoto algorithm to finite-state channels. *IEEE Transactions on Information Theory*, 54(5):1887–1918, 2008. 25

- [163] Tao Wang, Daniel Lizotte, Michael Bowling, and Dale Schuurmans. Bayesian sparse sampling for on-line reward optimization. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 956–963, 2005. [22](#)
- [164] Yizao Wang, Jean-Yves Audibert, and Rémi Munos. Algorithms for infinitely many-armed bandits. In *Proceedings of the 21st International Conference on Neural Information Processing Systems*, pages 1729–1736, 2008. [22](#)
- [165] Ward Whitt. Approximations of dynamic programs, I. *Mathematics of Operations Research*, 3(3):231–243, 1978. [23](#)
- [166] Hans S Witsenhausen. Separation of estimation and control for discrete time systems. *Proceedings of the IEEE*, 59(11):1557–1566, 1971. [22](#)
- [167] Yaming Yu. Squeezing the Arimoto–Blahut algorithm for faster convergence. *IEEE Transactions on Information Theory*, 56(7):3149–3157, 2010. [25](#)
- [168] Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pages 7304–7312. PMLR, 2019. [1](#), [22](#)
- [169] Julian Zimmert and Tor Lattimore. Connections between mirror descent, Thompson sampling and the information ratio. In *Advances in Neural Information Processing Systems*, pages 11973–11982, 2019. [7](#)
- [170] Luisa Zintgraf, Kyriacos Shiarlis, Maximilian Igl, Sebastian Schulze, Yarin Gal, Katja Hofmann, and Shimon Whiteson. VariBAD: A Very Good Method for Bayes-Adaptive Deep RL via Meta-Learning. In *International Conference on Learning Representations*, 2019. [25](#)

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#)
 - (c) Did you discuss any potential negative societal impacts of your work? [\[N/A\]](#)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[Yes\]](#)
 - (b) Did you include complete proofs of all theoretical results? [\[Yes\]](#)
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[N/A\]](#)
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[N/A\]](#)
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[N/A\]](#)
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[N/A\]](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[N/A\]](#)
 - (b) Did you mention the license of the assets? [\[N/A\]](#)
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[N/A\]](#)

- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]