

Exploration Potential

Jan Leike

*Future of Humanity Institute
University of Oxford*

JAN@FHI.AI

Editor:

Abstract

We introduce *exploration potential*, a quantity that measures how much a reinforcement learning agent has explored its environment class. In contrast to information gain, exploration potential takes the problem’s reward structure into account. This leads to an exploration criterion that is both necessary and sufficient for asymptotic optimality (learning to act optimally across the entire environment class). Our experiments in multi-armed bandits use exploration potential to illustrate how different algorithms make the tradeoff between exploration and exploitation.

1. Introduction

Good exploration strategies are currently a major obstacle for reinforcement learning (RL). The state of the art in deep RL (Mnih et al., 2015, 2016) relies on ϵ -greedy policies: in every time step, the agent takes a random action with some probability. Yet ϵ -greedy is a poor exploration strategy and for environments with sparse rewards it is quite ineffective (for example the Atari game ‘Montezuma’s Revenge’): it just takes too long until the agent randomwalks into the first reward.

More sophisticated exploration strategies have been proposed: using information gain about the environment (Sun et al., 2011; Orseau et al., 2013; Houthoof et al., 2016) or pseudo-count (Bellemare et al., 2016). In practice, these exploration strategies are employed by adding an exploration bonus (‘intrinsic motivation’) to the reward signal (Schmidhuber, 2010). While the methods above require the agent to have a model of its environment and formalize the strategy ‘explore by going to where the model has high uncertainty,’ there are also model-free strategies like the automatic discovery of options proposed by Machado and Bowling (2016). However, none of these explicit exploration strategies take the problem’s reward structure into account. Intuitively, we want to explore more in parts of the environment where the reward is high and less where it is low. This is readily exposed in optimistic policies like UCRL (Jaksch et al., 2010) and stochastic policies like PSRL (Strens, 2000), but these do not make the exploration/exploitation tradeoff explicitly.

In this paper, we propose *exploration potential*, a quantity that measures reward-directed exploration. We consider *model-based* reinforcement learning in partially or fully observable domains. Informally, exploration potential is the Bayes-expected absolute deviation of the value of optimal policies. Exploration potential is similar to information gain about the environment, but explicitly takes the problem’s reward structure into account. We show that this leads to a exploration criterion that is both *necessary* and *sufficient* for asymptotic optimality (learning to act optimally across an environment class): a reinforcement learning agent learns to act optimal in the limit *if and only if* the exploration potential converges to 0. As such, exploration potential captures the essence of what it means to ‘explore the right amount’.

Another exploration quantity that is both necessary and sufficient for asymptotic optimality is information gain about the optimal policy (Russo and Van Roy, 2014; Reddy et al., 2016). In contrast to exploration potential, it is not measured on the scale of rewards, making an explicit value-of-information tradeoff more difficult.

For example, consider a 3-armed Gaussian bandit problem with means 0.6, 0.5, and -1 . The information content is identical in every arm. Hence an exploration strategy based on maximizing information gain about the environment would query the third arm, which is easily identifiable as suboptimal, too frequently (linearly versus logarithmically). This arm provides information, but this information is not very useful for solving the reinforcement learning task. In contrast, an exploration potential based exploration strategy concentrates its exploration on the first two arms.

2. Preliminaries and Notation

A reinforcement learning agent interacts with an environment in cycles: at time step t the agent chooses an *action* a_t and receives a *percept* $e_t = (o_t, r_t)$ consisting of an *observation* o_t and a *reward* $r_t \in [0, 1]$; the cycle then repeats for $t + 1$. We use $\mathbf{x}_{<t}$ to denote a history of length $t - 1$. With abuse of notation, we treat histories both as outcomes and as random variables.

A *policy* is a function mapping a history $\mathbf{x}_{<t}$ and an action a to the probability $\pi(a \mid \mathbf{x}_{<t})$ of taking action a after seeing history $\mathbf{x}_{<t}$. An *environment* is a function mapping a history $\mathbf{x}_{1:t}$ to the probability $\nu(e_t \mid \mathbf{x}_{<t}a_t)$ of generating percept e_t after this history $\mathbf{x}_{<t}a_t$. A policy π and an environment ν generate a probability measure ν^π over infinite histories, the expectation over this measure is denoted with \mathbb{E}_ν^π . The *value* of a policy π in an environment ν given history $\mathbf{x}_{<t}$ is defined as

$$V_\nu^\pi(\mathbf{x}_{<t}) := (1 - \gamma) \mathbb{E}_\nu^\pi \left[\sum_{k=t}^{\infty} \gamma^k r_k \mid \mathbf{x}_{<t} \right],$$

where $\gamma \in (0, 1)$ is the discount factor. The *optimal value* is defined as $V_\nu^*(\mathbf{x}_{<t}) := \sup_\pi V_\nu^\pi(\mathbf{x}_{<t})$, and the *optimal policy* is $\pi_\nu^* := \arg \max_\pi V_\nu^\pi$. We use μ to denote the true environment.

We assume the nonparametric setting: let \mathcal{M} denote a countable class of environments containing the true environment μ . Let $w \in \Delta \mathcal{M}$ be a prior probability distribution on \mathcal{M} . After observing the history $\mathbf{x}_{<t}$ the prior w is updated to the posterior $w(\nu \mid \mathbf{x}_{<t}) := w(\nu)\nu(\mathbf{x}_{<t}) / (\sum_{\rho \in \mathcal{M}} w(\rho)\rho(\mathbf{x}_{<t}))$. A policy π is *asymptotically optimal in mean* iff for every $\mu \in \mathcal{M}$, $\mathbb{E}_\mu^\pi[V_\mu^*(\mathbf{x}_{<t}) - V_\mu^\pi(\mathbf{x}_{<t})] \rightarrow 0$ as $t \rightarrow \infty$.

3. Exploration Potential

We consider *model-based* reinforcement learning where the agent learns a model of its environment. With this model, we can estimate the value of any candidate policy. Concretely, let \hat{V}_t^π denote our estimate of the value of the policy π at time step t . We assume that the agent's learning algorithm satisfies *on-policy value convergence* (OPVC):

$$V_\mu^\pi(\mathbf{x}_{<t}) - \hat{V}_t^\pi(\mathbf{x}_{<t}) \rightarrow 0 \text{ as } t \rightarrow \infty \text{ } \mu^\pi\text{-almost surely.} \quad (1)$$

This does *not* imply that our model of the environment converges to the truth, only that we learn to predict the value of the policy that we are following. On-policy value convergence does not require that we learn to predict off-policy, i.e., the value of other policies. In particular, we might not learn to predict the value of the μ -optimal policy π_μ^* .

For example, a Bayesian mixture or an MDL-based estimator both satisfy OPVC if the true environment is the environment class; for more details, see Leike (2016, Sec. 4.2.3).

We define the \hat{V}_t -greedy policy as $\pi_{\hat{V}_t}^* := \arg \max_{\pi} \hat{V}_t^{\pi}$.

3.1 Definition

Definition 1 (Exploration Potential) *Let \mathcal{M} be a class of environments and let $\mathbf{x}_{<t}$ be a history. The exploration potential is defined as*

$$\text{EP}_{\mathcal{M}}(\mathbf{x}_{<t}) := \sum_{\nu \in \mathcal{M}} w(\nu \mid \mathbf{x}_{<t}) \left| V_{\nu}^{\pi_{\hat{V}_t}^*}(\mathbf{x}_{<t}) - \hat{V}_t^{\pi_{\hat{V}_t}^*}(\mathbf{x}_{<t}) \right|.$$

Intuitively, EP captures the amount of exploration that is still required before having learned the entire environment class. Asymptotically the posterior concentrates around environments that are compatible with the current environment. EP then quantifies how well the model \hat{V}_t understands the value of the compatible environments' optimal policies.

Remark 2 (Properties of EP)

- (i) $\text{EP}_{\mathcal{M}}$ depends neither on the true environment μ , nor on the agent's policy π .
- (ii) $\text{EP}_{\mathcal{M}}$ depends on the choice of the prior w and on the agent's model of the world \hat{V}_t .
- (iii) $0 \leq \text{EP}_{\mathcal{M}}(\mathbf{x}_{<t}) \leq 1$ for all histories $\mathbf{x}_{<t}$.

The last item follows from the fact that the posterior $w(\cdot \mid \mathbf{x}_{<t})$ and the value function V are bounded between 0 and 1.

3.2 Sufficiency

Proposition 3 (Bound on Optimality) *For all $\mu \in \mathcal{M}$,*

$$V_{\mu}^*(\mathbf{x}_{<t}) - V_{\mu}^{\pi_{\hat{V}_t}^*}(\mathbf{x}_{<t}) \leq \hat{V}_t^*(\mathbf{x}_{<t}) - V_{\mu}^{\pi_{\hat{V}_t}^*}(\mathbf{x}_{<t}) + \frac{\text{EP}_{\mathcal{M}}(\mathbf{x}_{<t})}{w(\mu \mid \mathbf{x}_{<t})}.$$

The bound of Proposition 3 is to be understood as follows.

$$\underbrace{V_{\mu}^*(\mathbf{x}_{<t}) - V_{\mu}^{\pi_{\hat{V}_t}^*}(\mathbf{x}_{<t})}_{\text{optimality of the greedy policy}} \leq \underbrace{\hat{V}_t^*(\mathbf{x}_{<t}) - V_{\mu}^{\pi_{\hat{V}_t}^*}(\mathbf{x}_{<t})}_{\text{OPVC}} + \underbrace{\text{EP}(\mathbf{x}_{<t})}_{\text{exploration potential}} / \underbrace{w(\mu \mid \mathbf{x}_{<t})}_{\text{posterior}}$$

If we switch to the greedy policy $\pi_{\hat{V}_t}^*$, then $\hat{V}_t^* - V_{\mu}^{\pi_{\hat{V}_t}^*} \rightarrow 0$ due to on-policy value convergence (1). This reflects how well the agent learned the environment's response to the Bayes-optimal policy. Generally, following the greedy policy does not yield enough exploration for EP to converge to 0. In order to get a policy π that is asymptotically optimal, we have to combine an exploration policy which ensures that $\text{EP} \rightarrow 0$ and then gradually phase out exploration by switching to the $\pi_{\hat{V}_t}^*$ -greedy policy. Because of property (i), the agent can compute its current EP value and thus check how close it is to 0. The higher the prior belief in the true environment μ , the smaller this value will be (in expectation).

3.3 Necessity

Definition 4 (Policy Convergence) *Let π and π' be two policies. We say the policy π converges to π' in μ^π -probability iff $|\hat{V}_t^\pi(\mathbf{x}_{<t}) - \hat{V}_t^{\pi'}(\mathbf{x}_{<t})| \rightarrow 0$ as $t \rightarrow \infty$ in \hat{V} .*

We assume that \hat{V}_t is continuous in the policy argument. If π converges to π' in total variation in the sense that $\pi(a | \mathbf{x}_{<k}) - \pi'(a | \mathbf{x}_{<k}) \rightarrow 0$ for all actions a and $k \geq t$, then π converges to π' in \hat{V} .

Definition 5 (Strongly Unique Optimal Policy) *An environment μ admits a strongly unique optimal policy iff there is a μ -optimal policy π_μ^* such that for all policies π if*

$$V_\mu^*(\mathbf{x}_{<t}) - V_\mu^\pi(\mathbf{x}_{<t}) \rightarrow 0 \text{ in } \mu^\pi\text{-probability,}$$

then π converges to π_μ^ in \hat{V} .*

Assuming that \hat{V}_t^π is continuous in π , an environment μ has a unique optimal policy if there are no ties in $\arg \max_a V_\mu^*(\mathbf{x}_{<t}a)$. Admitting a strongly unique optimal policy is an even stronger requirement because it requires that there exist no other policies that approach the optimal value asymptotically but take different actions (i.e., there is a constant gap in the argmax). For any finite-state (PO)MDP with a unique optimal policy that policy is also strongly unique.

Proposition 6 (Asymptotic Optimality \Rightarrow EP \rightarrow 0) *If the policy π is asymptotically optimal in mean in the environment class \mathcal{M} and each environment $\nu \in \mathcal{M}$ admits a strongly unique optimal policy, then $\text{EP}_{\mathcal{M}} \rightarrow 0$ in μ^π -probability for all $\mu \in \mathcal{M}$.*

If we don't require the condition on strongly unique optimal policies, then the policy π could be asymptotically optimal while $\text{EP} \not\rightarrow 0$: there might be another policy π' that is very different from any optimal policy π_μ^* , but whose μ -value approaches the optimal value: $V_\mu^* - V_\mu^{\pi'} \rightarrow 0$ as $t \rightarrow \infty$. Our policy π could converge to π' without EP converging to 0.

4. Exploration Potential in Multi-Armed Bandits

In this section we use experiments with multi-armed Bernoulli bandits to illustrate the properties of exploration potential. The class of Bernoulli bandits is $\Theta = [0, 1]^k$ (the arms' means). In each time step, the agent chooses an action (arm) $i \in \{1, \dots, k\}$ and receives a reward $r_t \sim \text{Bernoulli}(\theta_i^*)$ where $\theta^* \in \Theta$ is the true environment. Since Θ is uncountable, exploration potential is defined with an integral instead of a sum:

$$\text{EP}_\Theta(\mathbf{x}_{<t}) := \int_\Theta p(\theta | \mathbf{x}_{<t}) |\theta_{j(\theta)} - \hat{\theta}_{j(\theta)}| d\theta$$

where $p(\theta | \mathbf{x}_{<t})$ is the posterior distribution given the history $\mathbf{x}_{<t}$, $\hat{\theta} := \int_\Theta \theta p(\theta | \mathbf{x}_{<t}) d\theta$ is the Bayes-mean parameter, and $j(\theta) := \arg \max_i \theta_i$ is the index of the best arm according to θ .

Figure 1 shows the exploration potential of several bandit algorithms, illustrating how much each algorithm explores. Notably, optimally confident UCB (Lattimore, 2015) stops exploring around time step 700 and focuses on exploitation (because in contrast to the other algorithms it knows

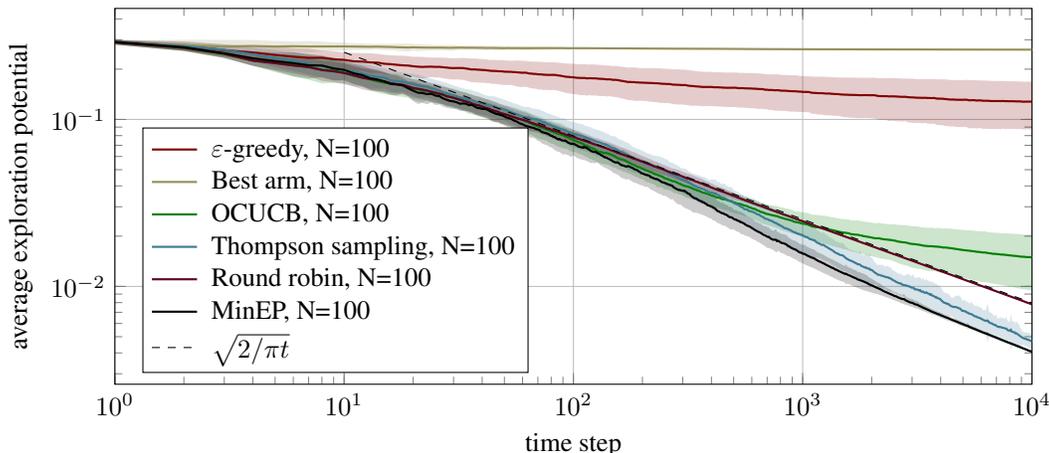


Figure 1: Exploration potential over time for different bandit algorithms in the Bernoulli bandit with arms 0.6, 0.5, 0.4, 0.4 (double logarithmic plot); shaded regions correspond to one standard deviation. Lower exploration potential means more exploration. The notable change in slope in around time steps 20–80 stems from the fact that it takes about that long to reliably distinguish the first two arms. The dashed line corresponds to the optimal asymptotic rate of $t^{-1/2}$.

the horizon). Thompson sampling, round robin (alternate between all arms), and ε -greedy explore continuously (but ε -greedy is less effective). The optimal strategy (always pull the first arm) never explores and hence its exploration potential decreases only slightly.

Exploration potential naturally gives rise to an exploration strategy: greedily minimize Bayes-expected exploration potential (MinEP); see Algorithm 1. This strategy unsurprisingly explores more than all the other algorithms when measured on exploration potential, but in bandits it also turns out to be a decent exploitation strategy because it focuses its attention on the most promising arms. For empirical performance see Figure 2. However, MinEP is generally not a good exploitation strategy in more complex environments like MDPs.

Algorithm 1 The MinEP Algorithm

- 1: **for** $t \in \mathbb{N}$ **do**
 - 2: $a_t := \arg \min_{a \in \mathcal{A}} \mathbb{E}_{e_t \sim \text{posterior}}[\text{EP}(a_{<t} a e_t)]$
 - 3: take action a_t
 - 4: observe percept e_t
-

5. Discussion

Several variants on the definition exploration potential given in Definition 1 are conceivable. However, often they do not satisfy at least one of the properties that make our definition appealing. Either they

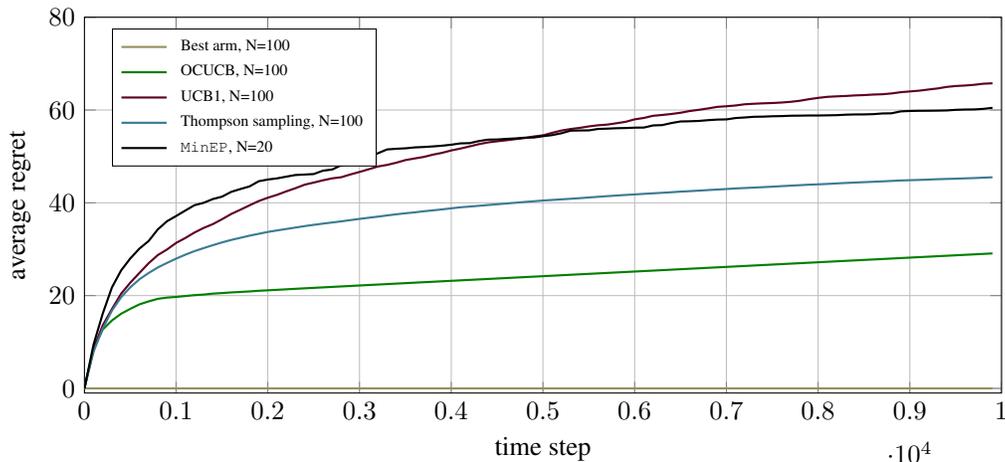


Figure 2: Average regret over time for different bandit algorithms in the Bernoulli bandit with arms 0.6, 0.5, 0.4, 0.4. MinEP outperforms UCB1 (Auer et al., 2002) after 10 000 steps, but neither Thompson sampling nor OCUCB.

break the necessity (Proposition 3), sufficiency (Proposition 6), our proofs thereof, or they make EP hard to compute. For example, we could replace $|V_\nu^* - \hat{V}_t^{\pi_\nu^*}|$ by $|V_\nu^* - V_\nu^\pi|$ where π is the agent’s future policy. This preserves both necessities and sufficiency, but relies on computing the agent’s future policy. If the agent uses exploration potential for taking actions (e.g., for targeted exploration), then this definition becomes a self-referential equation and might be very hard to solve. Following Dearden et al. (1999), we could consider $|V_\nu^* - \hat{V}_t^*|$ which has the convenient side-effect that it is model-free and therefore applies to more reinforcement learning algorithms. However, in this case the necessity guarantee (Proposition 6) requires the additional condition that the agent’s policy converges to the greedy policy $\pi_{\hat{V}_\nu^*}^*$. Moreover, this does not remove the dependence on a model since we still need a model class \mathcal{M} and a posterior.

Based on the recent successes in approximating information gain (Houthoofd et al., 2016), we are hopeful that exploration potential can also be approximated in practice. Since computing the posterior is too costly for complex reinforcement learning problems, we could (randomly) generate a few environments and estimate the sum in Definition 1 with them.

In this paper we only scratch the surface on exploration potential and leave many open questions. Is this the correct definition? What are good rates at which EP should converge to 0? Is minimizing EP the most efficient exploration strategy? Can we compute EP more efficiently than information gain?

References

Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.

Marc G Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. Technical report, 2016. <http://arxiv.org/abs/1606.01868>.

- Richard Dearden, Nir Friedman, and David Andre. Model based Bayesian exploration. In *Uncertainty in Artificial Intelligence*, pages 150–159, 1999.
- Rein Houthoofd, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. Curiosity-driven exploration in deep reinforcement learning via bayesian neural networks. Technical report, 2016. <http://arxiv.org/abs/1605.09674>.
- Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(Apr):1563–1600, 2010.
- Tor Lattimore. Optimally confident UCB: Improved regret for finite-armed bandits. Technical report, 2015. <http://arxiv.org/abs/1507.07880>.
- Jan Leike. *Nonparametric General Reinforcement Learning*. PhD thesis, Australian National University, 2016.
- Marlos C Machado and Michael Bowling. Learning purposeful behaviour in the absence of rewards. Technical report, 2016. <http://arxiv.org/abs/1605.07700>.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, 2016.
- Laurent Orseau, Tor Lattimore, and Marcus Hutter. Universal knowledge-seeking agents for stochastic environments. In *Algorithmic Learning Theory*, pages 158–172. Springer, 2013.
- Gautam Reddy, Antonio Celani, and Massimo Vergassola. Infomax strategies for an optimal balance between exploration and exploitation. *Journal of Statistical Physics*, 163(6):1454–1476, 2016.
- Dan Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. In *Advances in Neural Information Processing Systems*, pages 1583–1591, 2014.
- Jürgen Schmidhuber. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247, 2010.
- Malcolm Strens. A Bayesian framework for reinforcement learning. In *International Conference on Machine Learning*, pages 943–950, 2000.
- Yi Sun, Faustino Gomez, and Jürgen Schmidhuber. Planning to be surprised: Optimal bayesian exploration in dynamic environments. In *Artificial General Intelligence*, pages 41–51. Springer, 2011.