

Policy Optimization with Model-Based Explorations

Feiyang Pan,^{*1,2†} Qingpeng Cai,^{*3‡} An-Xiang Zeng,⁴ Chun-Xiang Pan,⁴ Qing Da,⁴
Hualin He,⁴ Qing He,^{1,2†} Pingzhong Tang^{3‡}

¹Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS),
Institute of Computing Technology, CAS, Beijing 100190, China. {panfeiyang, heqing}@ict.ac.cn

²University of Chinese Academy of Sciences, Beijing 100049, China.

³IIS, Tsinghua University. cqp14@mails.tsinghua.edu.cn, kenshinping@gmail.com

⁴Alibaba Group. {renzhong, xuanran}@taobao.com, {daqing.dq, hualin.hhl}@alibaba-inc.com

Abstract

Model-free reinforcement learning methods such as the Proximal Policy Optimization algorithm (PPO) have successfully applied in complex decision-making problems such as Atari games. However, these methods suffer from high variances and high sample complexity. On the other hand, model-based reinforcement learning methods that learn the transition dynamics are more sample efficient, but they often suffer from the bias of the transition estimation. How to make use of both model-based and model-free learning is a central problem in reinforcement learning.

In this paper, we present a new technique to address the trade-off between exploration and exploitation, which regards the difference between model-free and model-based estimations as a measure of exploration value. We apply this new technique to the PPO algorithm and arrive at a new policy optimization method, named Policy Optimization with Model-based Explorations (POME). POME uses two components to predict the actions' target values: a model-free one estimated by Monte-Carlo sampling and a model-based one which learns a transition model and predicts the value of the next state. POME adds the error of these two target estimations as the additional exploration value for each state-action pair, i.e., encourages the algorithm to explore the states with larger target errors which are hard to estimate. We compare POME with PPO on Atari 2600 games, and it shows that POME outperforms PPO on 33 games out of 49 games.

Introduction

Reinforcement Learning focuses on maximizing long-term return by interacting with the environment sequentially (Sutton and Barto 1998). Generally, a reinforcement learning algorithm has two aspects, on the one hand, it estimates the

state-action value function (also known as the Q-function), on the other hand, it optimizes or improves the policy to maximize its performance measure.

There are two classes of reinforcement learning methods, model-free and model-based methods. Model-free methods (Peters and Schaal 2006) estimate and iteratively update the state-action value with the rollout samples by Temporal Difference learning (Sutton and Barto 1998). It is said that model-based methods maintain an approximate model including the reward functions and the state transitions, and then use the approximated rewards and transitions to estimate the value function. Model-based methods are more efficient than model-free methods (Li and Todorov 2004; Levine and Koltun 2013; Montgomery and Levine 2016; Wahlström, Schön, and Desienroth 2015; Watter et al. 2015) especially in discrete environments by reducing the sample complexity. Model-free methods are more generally applicable to continuous and complex control problems but may suffer from high sample complexity (Schulman et al. 2015; Lillicrap et al. 2015).

Model-free methods directly use the immediate rewards and next states from rollout samples and estimate the long-term state-action value by Temporal Difference learning, for example, Sarsa or Q-learning (Sutton and Barto 1998). Therefore, the target value is unbiased but may induce large variance due to the randomness of the transition dynamics or the off-policy stochastic exploration strategy. Model-based methods use the prediction of the immediate reward and the next state by its own belief of the environment. The belief of the agent, including the approximate reward function and transition model, is updated after receiving more signals from the environment. So the estimated target value in model-based methods is often biased due to the approximation error of the model, but it has low variance compared with model-free methods. Combining model-based and model-free learning has been an important question in the recent literature.

We aim to answer the following question: How to incorporate model-based and model-free methods for better control?

The Dyna-Q (Sutton 1990), Normalized Advantage Function (NAF) algorithm (Gu et al. 2016), and Model-Ensemble Trust-Region Policy Optimization (ME-TRPO) algorithm (Kurutach et al. 2018) use simulated experiences in a learned

*These authors contributed equally in this work.

†The research was supported by the National Key Research and Development Program of China under Grant No. 2017YFB1002104, the National Natural Science Foundation of China under Grant No. 61602438, 91546122, 61573335, Guangdong provincial science and technology plan projects under Grant No. 2015 B010109005.

‡Pingzhong Tang and Qingpeng Cai were supported in part by the National Natural Science Foundation of China Grant 61561146398, a China Youth 1000-talent program and an Alibaba Innovative Research program.

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

transition model to supplement the real transition samples. (Cai, Pan, and Tang 2018) uses a convex combination of the two target values as the new target, following the insight that the ensemble could be more accurate. But the large loss of model prediction would also lead to an incorrect update for policy iteration.

In this paper, we incorporate the two methods together to address the trade-off between exploration and exploitation (EE), by a simple heuristic that adds the discrepancy between both target values as a relative measure of the *exploration value*, so as to encourage the agent to explore more difficult transition dynamics.

The trade-off between EE is a fundamental and challenging problem in reinforcement learning because it is hard to evaluate the value of exploring unfamiliar states and actions. Model-free exploration methods are widely discussed over decades. To name some, Upper Confidence Bounds (Auer, Cesa-Bianchi, and Fischer 2002; Li et al. 2010; Chen et al. 2017) adds the upper confidence bound of the uncertainty of value estimation as an exploration value of actions. Thompson Sampling (Thompson 1933) can also be understood as adding stochastic exploration value based on the posterior distribution of value estimation (May et al. 2012). However these model-free value-based methods often require a table to store the number of visits of each state-action pair so that the uncertainty of value estimation can be delivered (Tang et al. 2017), so they are not practical to handle problems with continuous state spaces such as Atari games or large-scale real-world problems such as mechanism design (Tang 2017) and e-commerce (Cai et al. 2018a; 2018b). Model-based methods for deep RL can have better exploration based on maximization of information gain (Houthoofd et al. 2016) or minimization of error (Pathak et al. 2017) about the agents’ belief, however, may be unstable due to the difficulty of learning the transition dynamics.

In this paper, based on the discrepancy between model-free and model-based target values, we present a new algorithm named Policy Optimization with Model-based Exploration (POME), an exploratory modified version of the well-known algorithm Proximal Policy Optimization (PPO). Different from previous methods, we use the model-free sample-based estimator for the advantage function as the base for POME, which is already successfully used in various actor-critic methods including PPO. Then POME adds a centralized and clipped exploration bonus onto the advantage value, so as to encourage exploration as well as stabilize the performance.

Our intuition is that the discrepancy of target values of model-based and model-free methods can be understood as the uncertainty of the transition dynamics. A high exploration value means that the transition of the state-action pair is hard to estimate and needs to be explored. We directly add the difference to the advantage estimation as the exploration bonus of the state-action pair. So if the exploration value of a state-action pair is higher than average, POME will encourage the agent to visit it more frequently in the future so as to better learn the transition dynamics. If not, POME will reduce the chance of picking it and give the agent an opportunity to try other actions.

We verify POME on the Atari 2600 game playing benchmarks. We compare our POME with the original model-free PPO algorithm and a model-based extended version of PPO. Experimental results show that POME outperforms the original PPO on 33 Atari games out of 49. We also tested two versions of POME, one with decaying exploration level and the other with non-decaying exploration level, which results in an interesting finding that the exploration bonus can improve the performance even in a long run.

Background and Notations

A Markov Decision Process (MDP) is defined as the tuple $(\mathcal{S}, \mathcal{A}, r, P, P_0)$, where \mathcal{S} is the (discrete or continuous) state space, \mathcal{A} is the (discrete or continuous) action space, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the (immediate) reward function, P is the state transition model, and P_0 is the probability distribution for the initial state s_0 .

The goal of the agent is to find the policy π^* from a restricted family of parametrized policy functions that maximizes its performance,

$$\pi^* = \arg \max_{\pi \in \Pi} J(\pi), \quad (1)$$

where $J(\pi)$ is the performance objective defined as

$$J(\pi) = \mathbb{E}_{s_0, a_0, s_1, \dots} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid \pi, s_0 \sim P_0 \right], \quad (2)$$

where $\gamma \in (0, 1)$ is a discount factor that balances the short- and long-term returns. For convenience, let $\rho^\pi(s)$ denote the (unnormalized) discounted cumulated state distribution induced by policy π ,

$$\rho^\pi(s) := \sum_{t=0}^{\infty} \gamma^t \Pr(s_t = s \mid \pi, s_0 \sim P_0) \quad (3)$$

then the performance objective can rewrite as

$$J(\pi) = \mathbb{E}_{s \sim \rho^\pi, a \sim \pi} [r(s, a)].$$

Since the expected long-term return is unknown, the basic idea behind RL is to construct a tractable estimator to approximate the actual return. Then we can update the policy in the direction that the performance measure is guaranteed to improve at every step (Kakade and Langford 2002). Let $Q^\pi(s, a)$ denote the expected action value of action a at state s by following policy π , i.e.

$$Q^\pi(s, a) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a, \pi \right]. \quad (4)$$

And we write the state value $V^\pi(s)$ and the advantage action value $A^\pi(s, a)$ as

$$V^\pi(s) = \int_{\mathcal{A}} \pi(a \mid s) Q^\pi(s, a) da, \quad (5)$$

$$A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s). \quad (6)$$

So the state value is defined as the weighted average of action values, and the advantage function provides a relative measure of value of each action.

Since we have the Bellman equation

$$Q^\pi(s, a) = r(s, a) + \gamma \mathbb{E}_{s'}[V^\pi(s')], \quad (7)$$

the advantage function can also be written as

$$A^\pi(s, a) = r(s, a) + \gamma \mathbb{E}_{s'}[V^\pi(s')] - V^\pi(s). \quad (8)$$

In practice, estimations of these quantities are used to evaluate the policy and to guide the direction of policy updates.

Policy Optimization Methods

Policy Gradient (PG) methods (Sutton et al. 2000) compute the gradient of $J(\pi_\theta)$ w.r.t policy parameters θ and then use gradient ascent to update the policy. The well-known form of policy gradient writes as follow

$$\nabla_\theta J(\pi_\theta) = \mathbb{E}_{s \sim \rho^{\pi_\theta}, a \sim \pi_\theta} [A^{\pi_\theta}(s, a) \nabla_\theta \log \pi_\theta(a|s)]. \quad (9)$$

Instead of directly using the policy gradients, trust region based methods use a surrogate objective to perform policy optimization. For the reason that the trajectory distribution is known only for the prior policy π_{old} before the update, trust region based methods introduce the following local approximation to $J(\pi_\theta)$ for any untested policy π_θ , as

$$L_{\pi_{\text{old}}}(\pi_\theta) = J(\pi_{\text{old}}) + \mathbb{E}_{s \sim \rho^{\pi_{\text{old}}}, a \sim \pi_\theta} [\pi_\theta(a|s) A^{\pi_{\text{old}}}(s, a)]. \quad (10)$$

It is proved (Kakade and Langford 2002; Schulman et al. 2015) that if the new policy π_θ is close to the π_{old} in the sense of the KL divergence, there is a lower bound of the long-term rewards of the new policy. For the general case, we have the following result:

Theorem 1 (Schulman et al. 2015). *Let $D_{\text{KL}}^{\text{max}}(\pi_{\text{old}}, \pi_\theta) = \max_s \text{KL}[\pi_{\text{old}}(\cdot|s) \|\pi_\theta(\cdot|s)]$ and $A_{\text{max}} = \max_s |\mathbb{E}_{a \sim \pi_\theta} [A^{\pi_{\text{old}}}(s, a)]|$. The performance of the policy $\tilde{\pi}$ can be bounded by*

$$J(\pi_\theta) \geq L_{\pi_{\text{old}}}(\pi_\theta) - \frac{2\gamma A_{\text{max}}}{(1-\gamma)^2} D_{\text{KL}}^{\text{max}}(\pi_{\text{old}}, \pi_\theta). \quad (11)$$

By following Theorem 1, there is a category of policy iteration algorithms with monotonic improvement guarantee, known as conservative policy iteration (Kakade and Langford 2002). Among them, Trust Region Policy Optimization (TRPO) is one of the most widely used baselines to optimize parametric policies to solve complex or continuous problems. The optimization problem w.r.t the new parameter θ is

$$\begin{aligned} \max_{\theta} \quad & \mathbb{E}_{s \sim \rho^{\pi_{\text{old}}}, a \sim \pi_{\text{old}}} \left[\frac{\pi_\theta(a|s)}{\pi_{\text{old}}(a|s)} A^{\pi_{\text{old}}}(s, a) \right] \\ \text{s.t.} \quad & \text{KL}[\pi_{\text{old}}(\cdot|s) \|\pi_\theta(\cdot|s)] \leq \delta_{\text{KL}} \text{ for all } s. \end{aligned} \quad (12)$$

where δ_{KL} is a hard constraint for the KL-divergence. In practice, the hard constraint of KL-divergence is softened by an expectation over visited states and moved to the objective with a multiplier β so that it becomes an unconstrained optimization problem, i.e.

$$\max_{\theta} \mathbb{E} \left[\frac{\pi_\theta(a|s)}{\pi_{\text{old}}(a|s)} A^{\pi_{\text{old}}}(s, a) - \beta \text{KL}[\pi_{\text{old}}(\cdot|s) \|\pi_\theta(\cdot|s)] \right], \quad (13)$$

which stabilise standard RL objectives.

Proximal Policy Optimization (PPO) (Schulman et al. 2017) can be view as a modified version of TRPO, which mainly uses a clipped probability ratio in the objective to avoid excessively large policy updates. It mainly replaces the term inside the expectation in (12) with

$$l^{\text{PPO}}(\theta; s, a) = \min \left\{ \frac{\pi_\theta(a|s)}{\pi_{\text{old}}(a|s)} A^{\pi_{\text{old}}}(s, a), \text{clip} \left(\frac{\pi_\theta(a|s)}{\pi_{\text{old}}(a|s)}, 1 - \epsilon, 1 + \epsilon \right) A^{\pi_{\text{old}}}(s, a) \right\}. \quad (14)$$

PPO is said to be empirically more sample efficient than TRPO, and is made one of the most frequently used baselines in various kinds of deep RL tasks.

Policy Evaluation with Function Approximation

In the previous part, we show several objectives for policy optimization. The objectives commonly need to compute an expectation over (s, a) pairs obtained by following the current policy and an estimator of advantage values.

In on-policy methods such as TRPO and PPO, the expectation is approximated with the empirical average over samples induced by the current policy π_{old} . Therefore the policy updates in a similar manner to the stochastic gradient descent as in modern supervised learning.

Model-free value estimation Model-free methods estimate the value functions by the samples in the rollout trajectory $(s_0, a_0, r_0, s_1, \dots)$. By following the Bellman equation, we can estimate the action or advantage value by Temporal Difference learning. When referring to the quantities at time step t of a trajectory, we will use Q_t , V_t , and A_t for short to denote $Q(s_t, a_t)$, $V(a_t)$, and $A(s_t, a_t)$ respectively.

Policy-based methods, as described in the previous part, typically need to use the advantage value A_t . It is discussed in (Schulman et al. 2015; Wang et al. 2016) that it can be one of the following estimators without introducing bias: the state-action value Q_t , the discounted return of the trajectory started from (s_t, a_t) , the one-step temporal difference error (TD error)

$$\delta_t = r_t + \gamma V_{t+1} - V_t, \quad (15)$$

or the k-step cumulated temporal difference error as used in the actual algorithm of PPO

$$\delta_t^k := \sum_{j=0}^{k-1} (\gamma \lambda)^j \delta_{t+j}, \quad (16)$$

where λ is a discount coefficient to balance future errors. When $\lambda = 1$, the estimator naturally becomes the advantage approximation in the asynchronous advantage actor critic algorithm (A3C) (Mnih et al. 2016) $A_{\text{A3C}}(s_t, a_t) = \sum_{j=0}^{k-1} \gamma^j r_{t+j} + \gamma^k V_{t+k} - V_t$. It is known that these estimators would not introduce bias to the policy gradient, but they have different variances.

Moreover, in deep RL, we often approximate the value functions with neural networks, which introduces additional biases and variances (Mnih et al. 2015; Schulman et al. 2015; Wang et al. 2016; Schulman et al. 2017). For example,

to compute the temporal difference errors, the state value $V^\pi(s)$ is always approximated with a neural network $V_\phi(s)$, where ϕ is the network parameter. Then if the one-step TD error is computed with on-policy samples, it becomes a loss function w.r.t the parameters ϕ ,

$$\delta_t(\phi) = Q_{f,t}^* - V_\phi(s_t), \quad (17)$$

where $Q_{f,t}^*$ denotes the model-free target value

$$Q_{f,t}^* = r_t + \gamma V_\phi(s_{t+1}). \quad (18)$$

Therefore in many algorithms the value parameters are updated by minimizing the error on samples using stochastic gradient descent.

The methods discussed so far can be categorized into model-free methods, because the immediate rewards and the state transitions are sampled from the trajectories.

Model-based value estimation For the model-based case, additional function approximation should be used to estimate the reward function and the transition function, i.e.

$$\hat{r}(s_t, a_t) \approx r_t \quad \text{and} \quad \hat{T}(s_t, a_t) \approx s_{t+1}. \quad (19)$$

In continuous problems the form of $\hat{r}(s_t, a_t)$ and $\hat{T}(s_t, a_t)$ can be neural networks, which are trained by minimizing some error measures between the model predictions and the real r_t and s_{t+1} . Here the symbols for parameters are omitted for simplicity. When using the mean squared error, we can define the loss function to these two networks

$$L_r = \mathbb{E}[r_t - \hat{r}(s_t, a_t)]^2, \quad L_T = \mathbb{E}\|s_{t+1} - \hat{T}(s_t, a_t)\|_2^2$$

which can be minimized using stochastic gradient descent w.r.t the network parameters. To make use of the model-based estimators, similar to the model-free case, we have the model-based TD error

$$\tilde{\delta}_t(\phi) = Q_{b,t}^* - V_\phi(s_t) \quad (20)$$

as an alternative to δ_t in (17), where $Q_{b,t}^*$ denotes the model-based target value

$$Q_{b,t}^* = \hat{r}(s_t, a_t) + \gamma V_\phi(\hat{T}(s_t, a_t)). \quad (21)$$

To solve complex tasks with high-dimensional inputs, model-free deep RL methods can often learn faster comparing to model-based methods especially in the beginning, mainly because it has fewer parameters to learn.

Policy Optimization with Model-based Explorations

In this section, we propose an extension of the trust region based policy optimization method with the exploration heuristics by making use of the difference between the model-free and model-based value estimations.

Let us think of the TD error as a subtract between the target value and the current function approximation. The target value for model-free and model-based learning is written in (18) and (21). Here we define the *discrepancy of targets* of state-action pair (s_t, a_t) as

$$\epsilon_t = |Q_{b,t}^* - Q_{f,t}^*| \quad (22)$$

$$= |r_t + \gamma V(s_{t+1}) - \hat{r}(s_t, a_t) - \gamma V(\hat{T}(s_t, a_t))|. \quad (23)$$

For our method, we use this error as an additional exploration value for the pair (s_t, a_t) . The proposed method is named Policy Optimization with Model-based Exploration (POME). So the resulted TD error used in POME is

$$\delta_t^{\text{POME}} = \delta_t + \alpha(\epsilon_t - \bar{\epsilon}), \quad (24)$$

where $\alpha > 0$ is a coefficient decaying to 0 and $\bar{\epsilon}$ is used to shift the exploration bonus to zero mean over samples from the same batch.

Insights

The basic idea behind POME is to add an exploration bonus to the state-action pairs where there is a discrepancy between the model-based and model-free Q-values.

The intuition is that the discrepancy between both Q-values would be small if the agent is “familiar” with the transition dynamics, for example, when the probability distribution of the next state concentrates on one deterministic state and at the meanwhile the next state has already been visited several times. On another hand, the error would be large if the agent is uncertain about what is going on, for example, if it has trouble predicting the next state or even it has never visited the next state yet. Therefore we think of the *discrepancy of targets* as a measure of the uncertainty of the long-term value estimation.

By the update rule of policy iteration methods (9) and (12), the chance that the policy selects the action with larger advantage estimation would be higher after a few updates. When starting to learn, the discrepancy of both target values can be high, which means that the transition dynamic is hard to estimate, thus we need to encourage the agent to explore. After training properly for a while, we tend to schedule α to a tiny value approaching 0, which means that the algorithm can asymptotically reach the model-free objective.

By the definition of (24), we let the discrepancy of targets serve as an additional exploration value of actions, along with the coefficient α to address the trade-off between exploration and exploitation. Note that, the exploration bonus on some state-action pairs will then be propagated to others so that our method guides exploration to these “interesting” parts of the state-action space.

Techniques

With the definition of the discrepancy of targets ϵ_t and the new exploratory target value δ_t^{POME} , we still need some techniques to build an actual policy optimization algorithm that is stable and efficient.

We first discuss why we use $\bar{\epsilon}$ for zero-mean normalization. Consider an on-policy algorithm that optimizes its policy using the samples $\{(s_t, a_t, r_t, s_{t+1})\}$ visited by following the old policy, i.e. $s \sim \rho^{\text{old}}, a \sim \pi_{\text{old}}$. Conventional methods estimate the advantage values on these state-action pairs to guide the next step of policy iteration. It means that, for a state s , if there exists a certain action a that has never been taken before, it has to wait for a moment that the algorithm gives negative advantage values to other actions at the same state, otherwise it will never get a chance of being chosen. However, when we use POME to give biased advantage

estimations by adding another term $\alpha(\epsilon_t - \bar{\epsilon})$ where $\bar{\epsilon}$ is used to normalize ϵ_t to zero mean, it naturally reduces the chance of exploiting familiar actions and encourages to pick unfamiliar actions which either have uncertain results or have little chance to be picked before. For the actual algorithm, $\bar{\epsilon}$ is calculated as the median of ϵ_t of on-policy samples because the median of samples is more robust to outliers than the mathematical average.

Next, we show a practical trick to stabilize the exploration. Since the discrepancy of targets ϵ_t is an unbounded positive value, it can be extremely large when the agent arrives at a totally unexpected state for the first time, so the algorithm would be unstable. So we clip the exploration value to a certain range to stabilize the algorithm. By following Theorem 1, it is necessary to reduce the error when estimating $\hat{A}(\cdot)$ to guarantee monotonic improvements of policy iteration. So we clip the discrepancy of targets to the same range of the model-free TD error, by replacing (24) with

$$\delta_t^{\text{POME}} = \delta_t + \alpha \text{clip}(\epsilon_t - \bar{\epsilon}, -|\delta_t|, |\delta_t|). \quad (25)$$

Algorithm 1 Policy Optimization with Model-based Explorations (single worker)

- 1: Input k // number of time steps per trajectory
 - 2: Initialize $w, \phi, \theta_{\hat{r}}, \theta_{\hat{T}}$ // parameters of π, V, \hat{r}, \hat{T}
 - 3: Initialize α // coefficient for exploration
 - 4: **while** not end **do**
 - 5: Sample a trajectories from the environment
 - 6: **for** $t = 1$ to k **do**
 - 7: $Q_{f,t}^* = r_t + \gamma V_{\phi}(s_{t+1})$
 - 8: $Q_{b,t}^* = \hat{r}(s_t, a_t) + \gamma V_{\phi}(\hat{T}(s_t, a_t))$
 - 9: $\epsilon_t = |Q_{b,t}^* - Q_{f,t}^*|$
 - 10: $\bar{\epsilon} = \text{Median}(\epsilon_1, \dots, \epsilon_{m=k})$
 - 11: **for** $t = k$ to 1 **do**
 - 12: $Q_{\text{POME},t}^* = Q_{f,t}^* + \alpha \text{clip}(\epsilon_t - \bar{\epsilon}, -|\delta_t|, |\delta_t|)$
 - 13: $\delta_t^{\text{POME}}(\phi) = Q_{\text{POME},t}^* - V_{\phi}(s_t)$
 - 14: **if** $t == m$ **then** $\hat{A}_t^{\text{POME}} = \delta_t^{\text{POME}}(\phi)$
 - 15: **else** $\hat{A}_t^{\text{POME}} = \delta_t^{\text{POME}}(\phi) + \hat{A}_{t+1}^{\text{POME}}$
 - 16: Update θ to maximize the surrogate objective (26)
 - 17: Update ϕ to minimize $L_v(\phi)$
 - 18: Update $\theta_{\hat{r}}, \theta_{\hat{T}}$ to minimize $L_r(\theta_{\hat{r}})$ and $L_T(\theta_{\hat{T}})$
 - 19: Decay the coefficient α towards 0
-

Formally, the optimization problem for POME is

$$\max_{\theta} \mathbb{E} \left[l^{\text{POME}}(\theta; s, a) - \beta \text{KL}[\pi_{\text{old}}(\cdot|s) \|\pi_{\theta}(\cdot|s)] \right], \quad (26)$$

where

$$l^{\text{POME}}(\theta; s, a) = \min \left\{ \frac{\pi_{\theta}(a|s)}{\pi_{\text{old}}(a|s)} \hat{A}^{\text{POME}}(s, a), \right. \\ \left. \text{clip} \left(\frac{\pi_{\theta}(a|s)}{\pi_{\text{old}}(a|s)}, 1 - \epsilon, 1 + \epsilon \right) \hat{A}^{\text{POME}}(s, a) \right\}. \quad (27)$$

Following the one-step temporal difference error defined in (25), we can use the k-step cumulated error to estimate the

advantage value of state-action pairs along the trajectory,

$$\hat{A}_t^{\text{POME}} := \delta_t^{k,\text{POME}} = \sum_{j \geq 0}^{k-1} (\gamma \lambda)^j \delta_{t+j}^{\text{POME}}. \quad (28)$$

Therefore it is easy to compute by simply replacing the advantage estimation. The overall sample-based objective for policy parameters θ is

$$L^{\text{POME}}(\theta) = \frac{1}{T} \sum_{t=1}^T \left[l^{\text{POME}}(\theta; s_t, a_t) - \beta \text{KL}[\pi_{\text{old}}(\cdot|s_t) \|\pi_{\theta}(\cdot|s_t)] \right]. \quad (29)$$

The loss function for value estimator therefore becomes

$$L_v(\phi) := \frac{1}{T} \sum_{t=1}^T [\hat{A}_t^{\text{POME}} + V_t - V_{\phi}(s_t)]^2. \quad (30)$$

Now we present our new algorithm named Policy Optimization with Model-based Exploration (POME) by incorporating the techniques above, as shown in Algorithm 1.

Experiments

Experimental Setup

We use the Arcade Learning Environment (Bellemare et al. 2013) benchmarks along with a standard open-sourced PPO implementation (Dhariwal et al. 2017).

We will evaluate two versions of POME in this section. For the first version, to guarantee that POME asymptotically approximates the original objective, we linearly anneal the coefficient α from 0.1 to 0 over the training period. For the second version, we fix the value of α as 0.1. The first version is used to show that the heuristic of POME helps fast learning because it mainly influences the algorithm in the beginning phase. The second version additionally shows that, even though the exploration value added to the estimation of advantage value introduces additional bias, the bias would not damage the performance too much.

We test on all 49 Atari games. Each game is run for 10 million timesteps, over 3 random seeds. The average score of the last 100 episodes will be taken as the measure of performance. The experimental results of PPO in the comparison of Table 1 and Table 2 are borrowed from the original paper (Schulman et al. 2017).

Implementation Details

For all the algorithms, the discount factor γ is set to 0.99 and the advantage values are estimated by the k-step error (16) where the horizon k is set to be 128. We use 8 actors (workers) to simultaneously run the algorithm and the mini-batch size is 128×8 .

For the basic network structures of PPO and POME, we use the same actor-critic architecture as (Mnih et al. 2016; Schulman et al. 2017), with shared convolutional layers and separate MLP layers for the policy and the value network.

Since POME is an extended version of PPO by incorporating model-based value estimations, the implementation of

POME involves two additional parts: fitting the model and adding the discrepancy of targets. Note that PPO and POME were always given the same amount of data during training.

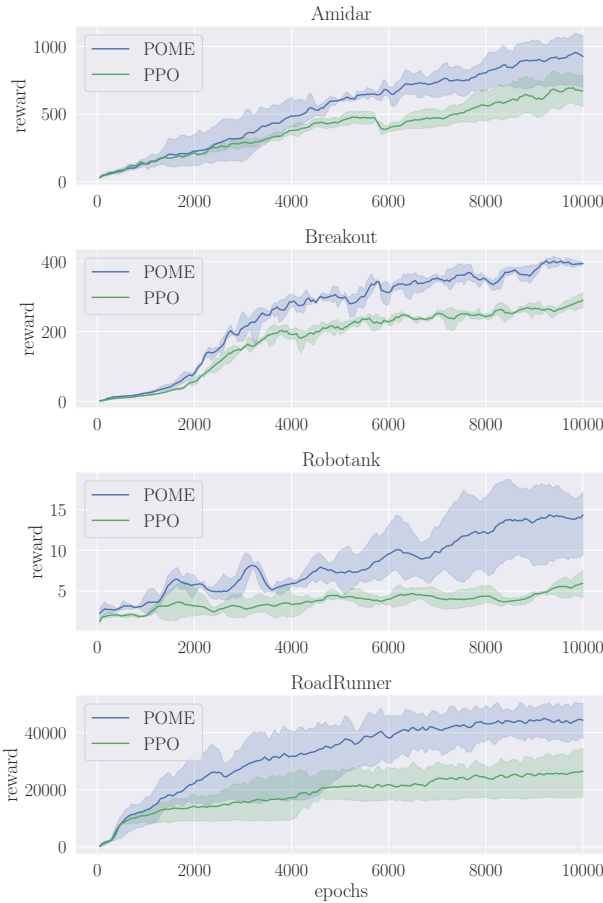


Figure 1: Comparison of POME against PPO on Atari games, training for 10M timesteps, over 3 random seeds.

To estimate the model, due to the fact that learning the state dynamics is more important than learning the rewards, we use a convolutional neural network with one hidden layer to fit the state transition. The inputs of the transition network have two parts, the state (images of four frames) and the action (a discrete number). Before being fed into the model, the images are scaled to the range $[0, 1]$, and the action is one-hot encoded. We concatenate the one-hot encoding of the action to the states' images to form the inputs of the transition model. We use the *sigmoid* activation for the outputs of the transition network. After that, we can finally compute the loss function of the transition model between the scaled images of the next state and the outputs.

In the actual implementation, POME use a unified objective function in order to simplify the computation

$$L = -L^{\text{POME}}(\theta) + c_v L_v(\phi) + c_T L_T(\theta_{\hat{T}}), \quad (31)$$

which is optimized by the Adam gradient descent optimizer (Kingma and Ba 2014) with learning rate $2.5 \times 10^{-4} \times f$, where f is a fraction linearly annealed from 1 to 0 over the

Table 1: Comparison between the original PPO and POME with decaying exploration coefficient. The scores of PPO are from the original paper (Schulman et al. 2017)

GAMES	PPO	POME
Alien	1850.3	1897.0
Amidar	674.6	943.9
Assault	4971.9	5638.6
Asterix	4532.5	4989.2
Asteroids	2097.5	1737.6
Atlantis	2311815.0	1941792.3
BankHeist	1280.6	1241.7
BattleZone	17366.7	15156.7
BeamRider	1590.0	1815.7
Bowling	40.1	58.3
Boxing	94.6	92.9
Breakout	274.8	411.8
Centipede	4386.4	2921.6
ChopperCommand	3516.3	4689.0
CrazyClimber	110202.0	115282.0
DemonAttack	11378.4	14847.1
DoubleDunk	-14.9	-6.8
Enduro	758.3	835.3
FishingDerby	17.8	21.1
Freeway	32.5	33.0
Frostbite	314.2	272.9
Gopher	2932.9	4801.8
Gravitar	737.2	914.5
IceHockey	-4.2	-4.5
Jamesbond	560.7	507.2
Kangaroo	9928.7	2511.0
Krull	7942.3	8001.1
KungFuMaster	23310.3	24570.3
MontezumaRevenge	42.0	0.0
MsPacman	2096.5	1966.5
NameThisGame	6254.9	5902.2
Pitfall	-32.9	-0.3
Pong	20.7	20.8
PrivateEye	69.5	100.0
Qbert	14293.3	15712.8
Riverraid	8393.6	8407.9
RoadRunner	25076.0	44520.0
Robotank	5.5	14.6
Seaquest	1204.5	1789.7
SpaceInvaders	942.5	964.2
StarGunner	32689.0	44696.7
Tennis	-14.8	-15.5
TimePilot	4232.0	4052.0
Tutankham	254.4	199.8
UpNDown	95445.0	181250.4
Venture	0	2.0
VideoPinball	37389.0	33388.0
WizardOfWor	4185.3	4301.7
Zaxxon	5008.7	6358.0

course of learning, and c_v, c_T are coefficients for tuning the learning rate of the value function and the transition function. In our experiments we set these coefficients to $c_v = 1$ and $c_T = 2$.

Comparison with PPO

Table 1 compares POME with decaying coefficient α against the original PPO, on the averaged scores of the last 100 episodes of algorithms with each environment. In Table 1, we see that, among the 49 games, POME with decaying coefficient α outperforms PPO in 32 games at the last 100 episodes.

The learning curves of four representative Atari games is shown in Figure 1. It shows that, in these environments, POME outperforms PPO over the entire training period, which indicates that it achieves fast learning and validate the power of our exploration technique by using the discrepancy of targets as exploration value.

Additional Experimental Results

We now investigate two questions: (1) how would POME perform if we do not tune the coefficient to 0? (2) how would the direct model-based extension of PPO perform?

For the first question, we set up the experiment to see if the exploration value used in POME would damage the performance in a long run. The coefficient α is now set to 0.1 for the entire training period. Secondly, we implement a model-based extension of PPO with the same architecture of the transition network as POME and replacing the target value with the model-based target value (21), so the agent can perform on-policy learning while maintaining the belief model.

We test the two extensions on Atari 2600 games. The setup of the environments and the hyper-parameters remain the same with the previous experiments.

The experimental results in Table 2 show that the model-based version is far from good. Only in one game can it outperforms the baseline. It shows that by using pure model-based PPO, the approximation errors introduced by fitting the model can substantially affect the performance.

However, POME with non-decaying coefficient turns out to be not only good but even better than POME with decaying coefficient. It outperforms the original PPO in 33 games out of 49. This result indicates that even though adding the exploration value would increase the bias when estimating the advantage functions, it is empirically not harmful to the policy optimization algorithms in most of the environments.

Conclusion and Discussion

Due to the challenge of the trade-off between exploration and exploitation in environments with continuous state space, in this paper, we propose a novel policy-based algorithm named POME, which uses the discrepancy between both target values of model-free and model-based methods to build a relative measure of the exploration value. POME uses several practical techniques to enable the exploration while being stable, i.e., the clipped and centralized exploration value. In the actual algorithm, POME builds on the model-free PPO algorithm and adds the exploration bonus to the estimation of the advantage function. Experiments show that POME outperforms the original PPO in 33 Atari games out of 49.

There is yet a limitation that, if the reward signal is extremely sparse, the error of the target values would be close

Table 2: Comparison among PPO, POME with constant exploration coefficient, and the model-based extension of PPO.

GAMES	PPO	PPO model-based	POME non-decay
Alien	1850.3	1386.4	1658.1
Amidar	674.6	27.7	704.0
Assault	4971.9	872.2	6211.5
Asterix	4532.5	1606.2	7235.0
Asteroids	2097.5	1456.8	1788.4
Atlantis	2311815	2864448	2030477
BankHeist	1280.6	159.4	1245.8
BattleZone	17366.7	2790.0	15313.3
BeamRider	1590.0	448.9	1989.2
Bowling	40.1	28.0	66.2
Boxing	94.6	52.5	92.6
Breakout	274.8	18.8	399.2
Centipede	4386.4	3343.6	2684.7
ChopperCommand	3516.3	1603.5	3886.3
CrazyClimber	110202.0	3640.0	112166.3
DemonAttack	11378.4	169.4	18877.9
DoubleDunk	-14.9	-17.9	-8.8
Enduro	758.3	92.2	862.8
FishingDerby	17.8	-62.7	19.3
Freeway	32.5	4.0	33.0
Frostbite	314.2	265.4	275.1
Gopher	2932.9	102.4	5050.4
Gravitar	737.2	150.0	773.8
IceHockey	-4.2	-4.5	-4.5
Jamesbond	560.7	37.2	871.8
Kangaroo	9928.7	1422.0	2237.0
Krull	7942.3	6091.0	8795.4
KungFuMaster	23310.3	12255.5	27667.0
MontezumaRevenge	42.0	0.0	0.0
MsPacman	2096.5	2059.7	2101.0
NameThisGame	6254.9	4485.1	5462.9
Pitfall	-32.9	-591.5	-4.4
Pong	20.7	16.1	20.8
PrivateEye	69.5	39.3	98.9
Qbert	14293.3	3308.2	14373.7
Riverraid	8393.6	3646.9	8358.3
RoadRunner	25076.0	4534.5	41740.7
Robotank	5.5	2.6	13.3
Seaquest	1204.5	592.4	1805.7
SpaceInvaders	942.5	355.4	971.6
StarGunner	32689.0	1497.5	33836.0
Tennis	-14.8	-22.6	-12.6
TimePilot	4232.0	3150.0	3970.7
Tutankham	254.4	99.0	159.6
UpNDown	95445.0	1589.4	203080.4
Venture	0.0	0.0	55.0
VideoPinball	37389.0	19720.4	31462.0
WizardOfWor	4185.3	1551.0	6050.7
Zaxxon	5008.7	1.5	6088.7

to 0. So POME would have little improvement in such situations. For the future work, it would be interesting to further extend the insight to address the exploration problem in the environments with sparse reward signals, for example, by incorporating with some reward-independent curiosity-based exploration methods.

References

- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* 47(2-3):235–256.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research* 47:253–279.
- Cai, Q.; Filos-Ratsikas, A.; Tang, P.; and Zhang, Y. 2018a. Reinforcement mechanism design for e-commerce. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*, 1339–1348.
- Cai, Q.; Filos-Ratsikas, A.; Tang, P.; and Zhang, Y. 2018b. Reinforcement mechanism design for fraudulent behaviour in e-commerce. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*.
- Cai, Q.; Pan, L.; and Tang, P. 2018. Generalized deterministic policy gradient algorithms. *arXiv preprint arXiv:1807.03708*.
- Chen, R. Y.; Sidor, S.; Abbeel, P.; and Schulman, J. 2017. UCB exploration via Q-ensembles. *arXiv preprint arXiv:1706.01502*.
- Dhariwal, P.; Hesse, C.; Klimov, O.; Nichol, A.; Plappert, M.; Radford, A.; Schulman, J.; Sidor, S.; Wu, Y.; and Zhokhov, P. 2017. Openai baselines. <https://github.com/openai/baselines>.
- Gu, S.; Lillicrap, T.; Sutskever, I.; and Levine, S. 2016. Continuous deep Q-learning with model-based acceleration. In *International Conference on Machine Learning*, 2829–2838.
- Houthoofd, R.; Chen, X.; Duan, Y.; Schulman, J.; De Turck, F.; and Abbeel, P. 2016. Vime: Variational information maximizing exploration. In *Advances in Neural Information Processing Systems*, 1109–1117.
- Kakade, S., and Langford, J. 2002. Approximately optimal approximate reinforcement learning. In *ICML*, volume 2, 267–274.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Kurutach, T.; Clavera, I.; Duan, Y.; Tamar, A.; and Abbeel, P. 2018. Model-ensemble trust-region policy optimization. In *International Conference on Learning Representations*.
- Levine, S., and Koltun, V. 2013. Guided policy search. In *International Conference on Machine Learning*, 1–9.
- Li, W., and Todorov, E. 2004. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *ICINCO (1)*, 222–229.
- Li, L.; Chu, W.; Langford, J.; and Schapire, R. E. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, 661–670.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- May, B. C.; Korda, N.; Lee, A.; and Leslie, D. S. 2012. Optimistic bayesian sampling in contextual-bandit problems. *Journal of Machine Learning Research* 13(Jun):2069–2106.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 1928–1937.
- Montgomery, W. H., and Levine, S. 2016. Guided policy search via approximate mirror descent. In *Advances in Neural Information Processing Systems*, 4008–4016.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning (ICML)*, volume 2017.
- Peters, J., and Schaal, S. 2006. Policy gradient methods for robotics. In *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, 2219–2225. IEEE.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International Conference on Machine Learning*, 1889–1897.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. In *International Conference on Learning Representations*.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.
- Sutton, R. S.; McAllester, D. A.; Singh, S. P.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, 1057–1063.
- Sutton, R. S. 1990. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Machine Learning Proceedings 1990*. Elsevier. 216–224.
- Tang, H.; Houthoofd, R.; Foote, D.; Stooke, A.; Chen, O. X.; Duan, Y.; Schulman, J.; DeTurck, F.; and Abbeel, P. 2017. # exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, 2753–2762.
- Tang, P. 2017. Reinforcement mechanism design. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 5146–5150.
- Thompson, W. R. 1933. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika* 25(3/4):285–294.
- Wahlström, N.; Schön, T. B.; and Desienroth, M. P. 2015. From pixels to torques: Policy learning with deep dynamical models. In *Deep Learning Workshop at the 32nd International Conference on Machine Learning*.
- Wang, Z.; Bapst, V.; Heess, N.; Mnih, V.; Munos, R.; Kavukcuoglu, K.; and de Freitas, N. 2016. Sample efficient actor-critic with experience replay. *arXiv preprint arXiv:1611.01224*.
- Watter, M.; Springenberg, J.; Boedecker, J.; and Riedmiller, M. 2015. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in neural information processing systems*, 2746–2754.