

NOVEL POLICY SEEKING WITH CONSTRAINED OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

In problem-solving, we humans tend to come up with different novel solutions to the same problem. However, conventional reinforcement learning algorithms ignore such a feat and only aim at producing a set of monotonous policies that maximize the cumulative reward. The resulting policies usually lack diversity and novelty. In this work, we aim at enabling the learning algorithms with the capacity of solving the task with multiple solutions through a practical novel policy generation workflow that can generate a set of diverse and well-performing policies. Specifically, we begin by introducing a new metric to evaluate the difference between policies. On top of this well-defined novelty metric, we propose to rethink the novelty-seeking problem through the lens of constrained optimization, to address the dilemma between the task performance and the behavioral novelty in existing multi-objective optimization approaches, we then propose a practical novel policy seeking algorithm, Interior Policy Differentiation (IPD), which is derived from the interior point method commonly known in the constrained optimization literature. Experimental comparisons on benchmark environments show IPD can achieve a substantial improvement over previous novelty-seeking methods in terms of both the novelty of generated policies and their performances in the primal task.¹.

1 INTRODUCTION

In the sense of learning through interactions with the environment, the scheme of reinforcement learning (RL) is conceptually similar to the emergence of intelligence (Sutton et al., 1998): an agent explores and exploits information of a given environment, learns to master some certain skills through trials and errors to gain as much reward as possible. When solving a problem, we humans could be creative to come up with multiple different solutions and gain insights from searching for diverse solutions. e.g., a self-containing example is the various approaches in RL research.

While the state-of-the-art algorithms have achieved superhuman performance in a variety of challenging tasks (Vinyals et al., 2019; Akkaya et al., 2019; Berner et al., 2019; Badia et al., 2020; Elbarbari et al., 2021), the task of encouraging individualized diversity ² of learned agents, on the other hand, is relatively under-explored. Different from conventional RL agents that are only learned through interactions with the external environment, novel policy generation is a task considering the differentiation among individual policies. The differentiation among policies can be explained as the social influence (Rogoff, 1990; Ryan & Deci, 2000; van Schaik & Burkart, 2011; Henrich, 2017; Harari, 2014) in social science literature. Although many works have been proposed applying social motivation to Multi-Agent Reinforcement Learning (MARL) settings (Jaques et al., 2019; Hughes et al., 2018; Sequeira et al., 2011; Peysakhovich & Lerer, 2017), how to motivate a single RL agent to perform differently against existing agents is still an open question.

In previous attempts for novel policy generation, there are three main challenges: (1) heuristically defined metric for novelty estimation is computational expensive (Zhang et al., 2019), (2) defining novelty reward for an entire episode yields additional challenge in credit assignment, and (3) solving

¹Code is included in the supplemental material.

²Note that this is different from diversity-driven exploration. While the latter focuses on the inner-diversity of a policy, in our work we focus on the inter-diversity between policies




Figure 1: The comparison of the standard policy gradient method without novelty seeking (left), multi-objective optimization method (middle), and our constrained optimization approach (right) for generating novel policies. The standard policy gradient method does not try actively to find novel solutions. The multi-objective optimization method may impede the learning procedure when the novelty gradient is being applied all the time, e.g., a random initialized policy will be penalized from getting closer to the previous policy due to the conflict of gradients, which limits the learning efficiency and the final performance. On the contrary, the novelty gradient of our constrained optimization approach will only be considered within a certain region to keep the policy being optimized away from highly similar solutions. Such an approach is more flexible and makes the multi-objective optimization method as its special case.

the problem under the formulation of multi-objective optimization leads to the performance decay in the original task. Fig. 1 compares the policy gradients of three cases, namely the one without novel policy seeking, novelty-seeking with multi-objective optimization, and novelty-seeking with constrained optimization methods, respectively.

In this work we take into consideration not only the novelty of a set of learned policies but also the performance of those novel policies in the primal task, when addressing the problem of novel-policy-generation. Our contributions can be summarized as follows:

1. We introduce a lightweight metric to compute the difference between policies with instant feedback at every timestep, to address the first two drawbacks of previous novel policy seeking methods discussed above;
2. We propose a constrained optimization formulation for novel policy generation and design a practical algorithm of IPD, resembling the interior point method in constrained optimization literature;
3. The proposed method, IPD, is evaluated on several continuous control benchmarks, showing the strengths of our constrained optimization solution for novelty-seeking can generate a series of diverse and well-performing policies, compared to previous multi-objective novel policy generation methods.

2 RELATED WORK

Intrinsic motivation methods. In previous work, different approaches are proposed to provide intrinsic motivation or intrinsic reward as a supplementary to the primal task reward for better exploration (Houthoofd et al., 2016; Pathak et al., 2017; Burda et al., 2018a;b; Liu et al., 2019). All those approaches use the weighted sum of two rewards, the primal rewards provided by environments and the intrinsic rewards provided by different heuristics. On the other hand, the work of DIAYN and DADS (Eysenbach et al., 2018; Sharma et al., 2019) learn diverse skills without extrinsic reward. Those approaches focus on decomposing diverse skills of a single policy, while our work focuses on learning diverse behaviors among a batch of policies for the same task.

Diverse policy generation methods. The work of Such et al. shows that different RL algorithms may converge to different policies for the same task (Such et al., 2018). On the contrary, we are interested in learning different policies through a single algorithm with the capability of avoiding local optimum. The work of Pugh et al. establishes a standard framework for understanding and comparing different approaches to search for quality diversity (QD) (Pugh et al., 2016). Conti et al. proposes a solution which avoids local optima as well as achieves higher performance by adding novelty search and QD to evolution strategies (Conti et al., 2018). The Task-Novely Bisector (TNB) (Zhang et al.,

2019) aims to solve novel policy generation problem by jointly optimize the extrinsic rewards and novelty rewards defined by an auto-encoder. In this work, we first adopt TNB in the constrained optimization framework, resulting in Constrained TNB, to demonstrate the dilemma between the task performance and novelty pursuance.

Constrained Markov Decision Process. The Constrained Markov Decision Process (CMDP) (Altman, 1999) considers the situation where an agent interacts with the environment under certain constraints. Formally, the CMDP can be defined as a tuple $(\mathcal{S}, \mathcal{A}, \gamma, r, c, C, P, s_0)$, where \mathcal{S} and \mathcal{A} are the state and action space; $\gamma \in [0, 1)$ is a discount factor; $r : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ and $c : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ denote the reward function and cost function; $C \in \mathbb{R}^+$ is the upper bound of permitted expected cumulative cost; $P(\cdot|s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ denotes the transition dynamics, and s_0 is the initial state. Denote the Markovian policy class as Π , where $\Pi = \{\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1], \sum_a \pi(a|\pi) = 1\}$. The learning objective of a policy for CMDP is to find a $\pi^* \in \Pi$, such that

$$\pi^* = \max_{\pi \in \Pi} \mathbb{E}_{\tau \sim \pi, s' \sim P} \left[\sum_{t=0}^{\infty} \gamma^t r(s, a, s') \right], \quad \text{s.t.} \quad \mathbb{E}_{\tau \sim \pi, s' \sim P} \left[\sum_{t=0}^{\infty} \gamma^t c(s, a, s') \right] \leq C, \quad (1)$$

where τ indicates a trajectory (s_0, a_0, s_1, \dots) and $\tau \sim \pi$ represents the distribution over trajectories following policy π : $a_t \sim \pi(\cdot|s_t)$, $s_{t+1} \sim P(\cdot|s_t, a_t)$; $t = 0, 1, 2, \dots$. Previous literature provide several approaches to solve CMDP (Achiam et al., 2017; Chow et al., 2018; Ray et al., 2019; Sun et al., 2021).

In Sec.2.1, we define a metric space that measures the difference between policies, which is the fundamental element for the proposed methods. In Sec.2.2, we develop a practical estimation method for this metric. Sec.2.3 describes the formulation of constrained optimization on novel policy generation. The implementations of two practical algorithms are further introduced in Sec.2.4.

We denote the policies as $\{\pi_{\theta_i}; \theta_i \in \Theta, i = 1, 2, \dots\}$, wherein θ_i represents parameters of the i -th policy, Θ denotes the whole parameter space. In this work, we focus on improving the behavioral diversity of policies from PPO (Schulman et al., 2017), thus we use Θ to represent Θ_{PPO} in this paper. It is worth noting that the proposed methods can be easily extended to other RL algorithms (Schulman et al., 2015; Lillicrap et al., 2015; Fujimoto et al., 2018; Haarnoja et al., 2018). To simplify the notation, we omit π and denote a policy π_{θ_i} as θ_i unless stated otherwise.

2.1 MEASURING THE DIFFERENCE BETWEEN POLICIES

We use the Wasserstein metric W_p (Rüschendorf, 1985; Villani, 2008; Arjovsky et al., 2017) to measure the distance between policies. Concretely, we consider the Gaussian-parameterized policies, where the W_p over two policies can be written in the closed form $W_2^2(\mathcal{N}(m_1, \Sigma_1), \mathcal{N}(m_2, \Sigma_2)) = \|m_1 - m_2\|^2 + \text{tr}[\Sigma_1 + \Sigma_2 - 2(\Sigma_1^{1/2}\Sigma_2\Sigma_1^{1/2})^{1/2}]$ as $p = 2$, where $m_1, \Sigma_1, m_2, \Sigma_2$ are mean and covariance metrics of the two normal distributions. In the following of this paper, we use D_W to denote the W_2 and it is worth noting that for the deterministic policy class, the trace term disappears and the only term involving the means remains, i.e., $D_W = \|m_1 - m_2\|$ for Dirac delta policies located at points m_1 and m_2 . This diversity metric satisfies the three properties of a metric, namely identity, symmetry as well as triangle inequality.

Proposition 1 (Metric Space $(\Theta, \overline{D}_W^q)$). *The expectation of $D_W(\cdot, \cdot)$ of two policies over any state distribution $q(s)$:*

$$\overline{D}_W^q(\theta_i, \theta_j) := \mathbb{E}_{s \sim q(s)} [D_W(\theta_i(a|s), \theta_j(a|s))], \quad (2)$$

is a metric on Θ , thus $(\Theta, \overline{D}_W^q)$ is a metric space.

The proof of Proposition 1 is straightforward. It is worth mentioning that although the KL-divergence is commonly applied in measuring differences between policies (Schulman et al., 2015; 2017; Hong et al., 2018). It can not be used directly as a metric for novelty quantification as the symmetric property do not hold. In this sense, the Jensen Shannon divergence D_{JS} and Total Variance Distance D_{TV} (Endres & Schindelin, 2003; Fuglede & Topsoe, 2004; Schulman et al., 2015) can be also used as alternative choices, we choose D_W in our work for that the Wasserstein metric better preserves the continuity (Arjovsky et al., 2017).

On top of the metric space $(\Theta, \overline{D}_W^q)$, the novelty of a policy can be computed as follows,

Definition 1 (Novelty of Policy). *Given a reference policy set Θ_{ref} such that $\Theta_{ref} = \{\theta_i^{ref}, i = 1, 2, \dots\}$, $\Theta_{ref} \subset \Theta$, the novelty $U(\theta|\Theta_{ref})$ of policy θ is the minimal difference between θ and all policies in the reference policy set, i.e.,*

$$U(\theta|\Theta_{ref}) := \min_{\theta_j \in \Theta_{ref}} \overline{D}_W^q(\theta, \theta_j), \quad (3)$$

Consequently, to encourage the discovery of novel policies discovery, typical novel policy generation methods tend to directly maximize the novelty of a new policy, i.e., $\max_{\theta} U(\theta|\Theta_{ref})$, where Θ_{ref} includes all existing policies.

2.2 ESTIMATION OF $\overline{D}_W^q(\theta_i, \theta_j)$ AND THE SELECTION OF $q(s)$

In practice, the calculation of $\overline{D}_W^q(\theta_i, \theta_j)$ is based on Monte Carlo estimation where we need to sample s from $q(s)$. Although in Eq.(2) $q(s)$ can be selected simply as a uniform distribution over the state space, there remains two obstacles: first, in a finite state space we can get precise estimation after establishing ergodicity, but problem arises when facing continuous state spaces due to the difficulty of efficiently obtaining enough samples; second, when s is sampled from a uniform distribution q , we can only get *sparse* episodic reward instead of *dense* online reward which is more useful in learning. Therefore, we make an approximation here based on importance sampling.

Formally, we denote the domain of $q(s)$ as $\mathcal{S}_q \subset \mathcal{S}$ and assume $q(s)$ to be a uniform distribution over \mathcal{S}_q , without loss of generality in later analysis. Notice \mathcal{S}_q is closely related to the algorithm being used in generating trajectories (Henderson et al., 2018). As we only care about the reachable regions of a certain algorithm (in this work, PPO), the domain \mathcal{S}_q can be decomposed by $\mathcal{S}_q = \lim_{N \rightarrow \infty} \bigcup_{i=1}^N \mathcal{S}_{\theta_i}$, where \mathcal{S}_{θ_i} denotes all the possible states a policy θ_i can visit given a starting state distribution.

In order to get online-reward, we estimate Eq.(2) with

$$\overline{D}_W^q(\theta_i, \theta_j) = \mathbb{E}_{s \sim q(s)} [D_W(\theta_i(a|s), \theta_j(a|s))] = \mathbb{E}_{s \sim \rho_{\theta_i}(s)} \left[\frac{q(s)}{\rho_{\theta_i}(s)} D_W(\theta_i(a|s), \theta_j(a|s)) \right], \quad (4)$$

where we use $\rho_{\theta}(s)$ to denote the stationary state visitation frequency under policy θ , i.e., $\rho_{\theta}(s) = P(s_0 = s|\theta) + P(s_1 = s|\theta) + \dots + P(s_T = s|\theta)$ in finite horizon problems. We propose to use the averaged stationary visitation frequency as $q(s)$, e.g., for PPO, $q(s) = \bar{\rho}(s) = \mathbb{E}_{\theta \sim \Theta_{PPO}} [\rho_{\theta}(s)]$. Clearly, choosing $q(s) = \bar{\rho}(s)$ will be much better than choosing a uniform distribution as the importance weight will be closer to 1. Such an importance sampling process requires a necessary condition that $\rho_{\theta_i}(s)$ and $q(s)$ have the same domain, which can be guaranteed by applying a sufficient exploration noise on θ .

Another difficulty lies in the estimation of $\bar{\rho}(s)$, which is always intractable given a limited number of trajectories. However, during training, θ_i is a policy to be optimized and $\theta_j \in \Theta_{ref}$ is a fixed reference policy. The error introduced by approximating the importance weight as 1 will get larger when θ_i becomes more distinct from normal policies, at least in terms of the state visitation frequency. We may just regard increasing of the approximation error as the discovery of novel policies.

Proposition 2 (Unbiased Single Trajectory Estimation). *The estimation of $\rho_{\theta}(s)$ using a single trajectory τ is unbiased.*

The Proposition 2 follows the usual trick in RL that uses a single trajectory to estimate the stationary state visitation frequency, and indicates that with a sufficiently large number of trajectories, we are able to access the exact value of novelty. Given the definition of novelty and a practically unbiased sampling method, the next step is to develop an efficient learning algorithm.

2.3 CONSTRAINED OPTIMIZATION FORMULATION FOR NOVEL POLICY GENERATION

In the conventional RL paradigm, maximizing the expectation of cumulative rewards is commonly used as the objective. i.e., $\max_{\theta \in \Theta} \mathbb{E}_{\tau \sim \theta} [g]$, where $g = \sum_{t=0}^{\infty} \gamma^t r_t$ and $\tau \sim \theta$ denotes a trajectory τ sampled from the policy θ .

To improve the diversity of different agents' behaviors, the learning objective must take both the reward from the primal task and the policy novelty into consideration. Previous approaches (Houthooft

et al., 2016; Pathak et al., 2017; Burda et al., 2018a;b; Liu et al., 2019) often directly use the weighted sum of these two terms as the objective:

$$\max_{\theta \in \Theta} \mathbb{E}_{\tau \sim \theta} [g_{\text{total}}] = \max_{\theta \in \Theta} \mathbb{E}_{\tau \sim \theta} [\alpha \cdot g_{\text{task}} + (1 - \alpha) \cdot g_{\text{int}}], \quad (5)$$

where $0 < \alpha < 1$ is a weight hyper-parameter, g_{task} is the reward from the primary task, and $g_{\text{int}} = \sum_{t=0} \gamma^t r_{\text{int},t}$ is the cumulative intrinsic reward of the *intrinsic reward* $r_{\text{int},t}$. In our case, the intrinsic reward is the novelty reward $r_{\text{int}} = \min_{\theta_j \in \Theta_{\text{ref}}} \bar{D}_W^\rho(\theta, \theta_j)$. These methods can be summarized as Weighted Sum Reward (WSR) methods. Such an objective is sensitive to the selection of α as well as the formulation of r_{int} . For example, in our case formulating the novelty reward r_{int} as $\min_{\theta_j} \bar{D}_W^\rho(\theta, \theta_j)$, $\exp[\min_{\theta_j} \bar{D}_W^\rho(\theta, \theta_j)]$ and $-\exp[-\min_{\theta_j} \bar{D}_W^\rho(\theta, \theta_j)]$ will lead to significantly different results as they determine the trade-offs in the two terms given α . Besides, dilemma also arises in the selection of α : while a large α may undermine the contribution of intrinsic reward, a small α could ignore the importance of the primal task, leading to the failure of an agent in solving the task.

The crux of tackling such an issue is to deal with the conflict between different objectives. The work of Zhang et al. proposes the TNB, where the task reward is regarded as the dominant one while the novelty reward is regarded as subordinate. However, as TNB considers the novelty gradient all the time, it may hinder the learning process. Intuitively, well-performing policies should be more similar to each other than to random initialized policies. As a new random initialized policy is different enough from previous policies, considering the novelty gradient at beginning of training will result in a much slower learning process. We also observe such a problem in experiments that TNB not only fails to find well-performing policies, but also fails in the novelty seeking as all policies are driven away from original performance-pursuing-only policies.

In order to tackle the above problems and adjust the extent of novelty in new policies, we propose to solve the novelty-seeking problem under the perspective of constrained optimization. The intuition is as follows: while the task reward is considered as a learning objective, the novelty reward should be considered as a bonus instead of another objective, thus should not impede the learning of the primal task. Fig. 1 illustrates how novelty gradients impede the learning of a policy: at the beginning of learning, a random initialized policy should learn to be more similar to a well-performing policy rather than be different. The seeking of novelty should not be taken into consideration all the time during learning. With such an insight, we update the multi-objective optimization problem in Eq.(5) into a constrained optimization problem as:

$$\max_{\theta \in \Theta} f(\theta) = \mathbb{E}_{\tau \sim \theta} [g_{\text{task}}], \quad \text{s.t.} \quad g_t(\theta) = \bar{r}_{\text{int},t} - r_0 \geq 0, \forall t = 1, 2, \dots, T, \quad (6)$$

where r_0 is a threshold indicating minimal permitted novelty. In practice, we can use the averaged diversity of a batch of agents trained with PPO as r_0 , motivating our diversity-seeking policies to achieve a higher novelty, therefore, there is only a little need for a hyper-parameter tuning. $\bar{r}_{\text{int},t}$ denotes a moving average of $r_{\text{int},t}$ as we need not force every single action of a new agent to be different from others. Instead, we care more about the long-term differences. Therefore, we use cumulative novelty terms as constraints. Moreover, the constraints can be flexibly applied after the first t_S timesteps (e.g., $t_S = 20$) for the consideration of similar starting sequences, so that the constraints can be written as $g_t(\theta) \geq 0, \forall t = t_S, \dots, T$.

2.4 PRACTICAL NOVEL POLICY GENERATION METHODS

One thing to note is that WSR and TNB proposed in the prior work (Zhang et al., 2019) correspond to different approaches in constrained optimization problems, yet some important ingredients are missing. We in this section adopt TNB according to the Feasible Direction Method in constrained optimization and then propose our method of Interior Policy Differentiation (IPD), according to the Interior Point Method in constrained optimization literature. A detailed discussion on WSR is provided in Appendix D.

TNB: Feasible Direction Method The Feasible Direction Method (FDM) (Ruszczynski, 1980; Herskovits, 1998) solves the constrained optimization problem by finding a direction \vec{p} where taking gradient upon will lead to the increment of the objective function as well as constraints satisfaction,

Algorithm 1 IPD (Ours)

```

1: Input: (1) a behavior policy  $\theta_{old}$ , (2) a set of previous policies  $\{\theta_j\}_{j=1}^M$ , (3) a novelty metric  $U(\theta, \{\theta_j\}|\rho) = U(\theta, \{\theta_j\}|\tau) = \min_{\theta_j} \bar{D}_W^r(\theta, \theta_j)$ , (4) a novelty threshold  $r_0$  and starting point  $t_S$ .
2: Initialize  $\theta_{old}$ .
3: for iteration = 1, 2, ... do
4:   for t = 1, 2, ..., T do
5:     Step the environment by taking action  $a_t \sim \theta_{old}$  and collect transitions.
6:     if  $U(\theta_{old}, \{\theta_j\}|\tau) - r_0 < 0$  AND  $t > t_S$  then
7:       Break this episode.
8:     end if
9:   end for
10:  Update policy parameters based on sampled data.
11: end for

```

Algorithm 2 Constrained TNB (For Comparison)

```

1: Input: (1-4) same as IPD, (5) a value network for cost  $V_c$ .
2: Initialize  $\theta_{old}$ .
3: for iteration = 1, 2, ... do
4:   for t = 1, 2, ..., T do
5:     Step the environment by taking action  $a_t \sim \theta_{old}$  and collect transitions.
6:   end for
7:   Compute advantage of reward  $\hat{A}_{r,1}, \dots, \hat{A}_{r,T}$ .
8:   Compute advantage of cost  $\hat{A}_{c,1}, \dots, \hat{A}_{c,T}$ .
9:   Optimize objective for reward  $\mathcal{L}_r^{\text{CLIP}}$ , with gradient  $g_r = \nabla_{\theta} \mathcal{L}_r^{\text{CLIP}}$ .
10:  Optimize objective for cost  $\mathcal{L}_c^{\text{CLIP}}$ , with gradient  $g_c = -\nabla_{\theta} \mathcal{L}_c^{\text{CLIP}}$ .
11:  if  $U(\theta_{old}, \{\theta_j\}|\tau) - r_0 < 0$  then
12:    Calculate  $\vec{p}$  according to Eq.(7) with  $g_r$  and  $g_c$ 
13:  else
14:    Calculate  $\vec{p}$  with  $g_r$ 
15:  end if
16:  Update policy parameters
17: end for

```

i.e., $\nabla_{\theta} f^T \cdot \vec{p} > 0$, if $g > 0$ and $\nabla_{\theta} g^T \cdot \vec{p} > 0$ otherwise. The TNB proposes to use a revised bisector of gradients $\nabla_{\theta} f$ and $\nabla_{\theta} g$ as \vec{p} ,

$$\vec{p} = \begin{cases} \nabla_{\theta} f + \frac{|\nabla_{\theta} f|}{|\nabla_{\theta} g|} \nabla_{\theta} g, & \text{if } \cos(\nabla_{\theta} f, \nabla_{\theta} g) > 0 \\ \nabla_{\theta} f + \frac{|\nabla_{\theta} f|}{|\nabla_{\theta} g|} \nabla_{\theta} g \cdot \cos(\nabla_{\theta} f, \nabla_{\theta} g), & \text{otherwise,} \end{cases} \quad (7)$$

Clearly, Eq.(7) satisfies the constraints but it is more strict than it as the $\nabla_{\theta} g$ term always exists during the optimization of TNB. Based on TNB, we provide a revised approach, named Constrained Task Novel Bisector (CTNB), which resembles better with FDM. Specifically, when $g > 0$, CTNB will not apply $\nabla_{\theta} g$ on g . It is clear that TNB is a special case of CTNB when the novelty threshold r_0 is set to infinity. We note that in both TNB and CTNB, the learning stride is fixed to be $\frac{|\nabla_{\theta} f| + |\nabla_{\theta} g|}{2}$ and may lead to problem when $\nabla_{\theta} f \rightarrow 0$, where the final optimization result will rely heavily on the selection of g , i.e., the shape of g is crucial for the success of this approach. We propose CTNB in our work, as a constrained optimization variant of TNB, to demonstrate the importance of the constrained optimization perspective in novelty seeking, however, we do not in practice observe such a method achieves satisfactory performance.

IPD: Interior Point Method The Interior Point Method (Potra & Wright, 2000; Dantzig & Thapa, 2006) is another approach used to solve the constrained optimization problem. Thus here we solve Eq.(6) using the Interior Policy Differentiation (IPD), which can be regarded as an analogy of the Interior Point Method. In the vanilla Interior Point Method, the constrained optimization problem in Eq.(6) is solved by reforming it to an unconstrained form with an additional barrier term $-\alpha \log g(\theta)$ in the objective as $\max_{\theta \in \Theta} f(\theta) - \alpha \log g(\theta)$, or more precisely in our problem with the formulation with Eq.(6) we have $\max_{\theta \in \Theta} \mathbb{E}_{\tau \sim \theta} [g_{\text{task}} - \sum_{t=0}^T \alpha \log(\bar{r}_{\text{int},t} - r_0)]$, where $\alpha > 0$ is the barrier factor. Besides the log barrier term, there are other choices like $\alpha \frac{1}{g(\theta)}$ can be used and the objective becomes $\max_{\theta \in \Theta} f(\theta) + \alpha \frac{1}{g(\theta)}$. As α is small, the barrier term will introduce only minuscule influence on the objective. On the other hand, when θ get closer to the barrier, the objective will increase rapidly. The limits when $\alpha \rightarrow 0$ then lead to the solution of Eq.(6). The convergence of such methods are provided in previous works (Conn et al., 1997; Wright, 2001).

However, directly applying IPM is computationally expensive and numerically unstable. In this work, we propose a simple yet novel heuristic method that resembles the idea of barrier methods: we implicitly apply such barrier terms by providing termination signals in interactions with the environments. Our method can be regarded as revising the primal task MDP into a new one in which the behaviors of agents must satisfy novelty constraints. Specifically, in the RL paradigm, the learning procedure of an agent is determined by the experiences collected during interactions with the




Figure 2: Experimental results on the Four Reward Maze Problem. We generate 5 policies with different novel policy generation methods. In each figure, 5 trajectories start from the +1 reward point. Thresholds in CTNB and IPD are set as the averaged novelty of PPO policies.

environment and the sampling strategy used to filter experiences in the calculation of policy gradients. Since the learning process is based on sampled transitions, a more natural way can thus be used to perform the constrained optimization. We can simply bound the collected transitions in the feasible region by permitting previously trained M policies $\theta_i \in \Theta_{\text{ref}}, i = 1, 2, \dots, M$ sending termination signals during the training process of new agents. In other words, we implicitly bound the feasible region by terminating any new agent that steps outside it.

Consequently, during the training process, all valid samples we collected are inside the feasible region, which means these samples are less likely to appear in previously trained policies. At the end of the training, we obtain a new policy that has sufficient novelty. In this way, we no longer need to consider the trade-off between intrinsic and extrinsic rewards deliberately. The learning process of IPD is thus more robust and no longer suffers from the objective inconsistency.

Remark 1 (Reward-Shaping-Agnostic Novelty Seeking). *IPD is a gradient-free method with regard to the novelty reward.*

Remark 1 is an important property that only IPD owns. For other approaches, including both multi-objective approaches and constrained optimization approaches, an elaborated design of the novelty reward function is needed, e.g., it can be any monotonic increasing function of \bar{D}_W^{ρ} . While it is well-known that reward shaping (Randlöv & Alström, 1998; Laud, 2004) is in general a non-trivial work that requires domain knowledge (Vinyals et al., 2019; Akkaya et al., 2019; Berner et al., 2019; Elbarbari et al., 2021), such an additional novelty-reward term further increases the burden of proper reward design.

Differently, the seeking of novelty in IPD does not require any (policy) gradient information that flows from the novelty reward, therefore, the selection of reward scaling function is agnostic to the performance of IPD. IPD learns to become novel in a passive manner, i.e., an episode will be terminated whenever the averaged step-wise novelty is lower than a given threshold. Searching in a constant hyper-parameter space is at least tractable while searching in a monotonically increasing functional reward shaping class (Zhang et al., 2019) is not – let alone IPD works well with a default novelty threshold parameter $r_0 = \text{averaged differences between PPO policies}$.

3 EXPERIMENTS

According to Proposition 2, the novelty reward r_{int} in Eq.(6) under our novelty metric can be unbiasedly approximated by $r_{\text{int}} = \min_{\theta_j \in \Theta_{\text{ref}}} \bar{D}_W^{\rho}(\theta(a|s_t), \theta_j(a_j|s_t))$. We thus utilize this novelty metric directly throughout our experiments. We apply different novel policy generation methods, namely WSR, TNB, CTNB, and IPD, to the backbone RL algorithm PPO (Schulman et al., 2017). The extension to other popular RL algorithms is straightforward. More implementation details are depicted in Appendix E.

Experiments in the work of Henderson et al. (2018) show that one can simply change the random seeds before training to get policies that perform differently. Therefore, we use PPO with varying random seeds as a baseline method for novel policy generation and use the averaged differences between policies learned by this baseline as the **default threshold** r_0 in CTNB and IPD. Algorithm 1 and Algorithm 2 show the pseudo code of IPD and CTNB based on PPO, where the blue lines show the additional code added to the standard PPO. Qualitative results can be found in Appendix G.




Figure 3: The performance and novelty comparison of different methods in Hopper-v3, Walker2d-v3 and HalfCheetah-v3 environments. The value of novelty is normalized to relative novelty by regarding the averaged novelty of PPO policies as the baseline. The results are from 10 policies of each method, with the points showing their mean and lines showing their standard deviation.

3.1 THE FOUR REWARD MAZE PROBLEM

We first utilize a basic 2-D environment named Four Reward Maze as a diagnostic environment where we can visualize learned policies directly. In this environment, four positive rewards of different values (e.g., +5, +5, +10, +1 for top, down, left and right respectively) are assigned to four middle points with radius 1 on each edge in a 2-D $N \times N$ square map. We use $N = 16$ in our experiments. The observation of a policy is the current position and the agent will receive a negative reward of -0.01 at each timestep except stepping into the reward regions. Each episode starts from a randomly initialized position and the action space is limited to $[-1, 1]$. The performance of each agent is evaluated by the averaged performances over 100 trials.

Results are shown in Fig. 2, where the behaviors of the PPO agents are quite similar, suggesting the diversity provided by random seeds is limited. WSR and TNB solve the novel policy generation problem from the multi-objective optimization formulation, they thus suffer from the unbalance between performance and novelty. While WSR and TNB both provide sufficient novelty, performances of agents learned by WSR decay significantly, so did TNB due to an encumbered learning process, as we analyzed in Sec.2.3. Both CTNB and IPD, solving the task with novelty-seeking from the constrained optimization formulation, provide evident behavior diversity and perform recognizably better than TNB and WSR.

3.2 THE MUJoCo BENCHMARK

We evaluate our proposed method on three locomotion tasks (Brockman et al., 2016; Todorov et al., 2012): the Hopper-v3 (11 observations and 3 actions), Walker2d-v3 (11 observations and 6 actions), and HalfCheetah-v3 (17 observations and 6 actions). Although relaxing the healthy termination thresholds in Hopper and Walker may permit more visible behavior diversity, all the environment parameters are set as default values in our experiments to demonstrate the generality of our method.

3.2.1 COMPARISON ON NOVELTY AND PERFORMANCE

We implement WSR, TNB, CTNB, and IPD using the same hyper-parameter settings per environment. And we also apply CPO (Achiam et al., 2017) as a baseline as a solution of CMDP. For each method, we first train 10 policies using PPO with different random seeds. Those PPO policies are used as the primal reference policies, and then we train 10 novel policies that try to be different from previous reference policies. Concretely, in each method, the 1st novel policy is trained to be different from the previous 10 PPO policies, and the 2nd should be different from the previous 11 policies, and so on. More implementation details are depicted in Appendix E.

Fig. 3 shows the results in terms of novelty (the x-axis) and the performance (the y-axis). Policies close to the upper right corner are the more novel ones with higher performance. In all environments, the performance of CTNB, IPD and CPO outperforms WSR and TNB, showing the advantage of constrained optimization approaches in novel policy generation. Specifically, the results of CTNB are

Table 1: Reward and Success Rate of 10 Policies. IPD beat CTNB, CPO, TNB and WSR in all three environments. Constrained optimization approaches outperforms multi-objective methods.

Environment	Reward			Success Rate		
	Hopper	Walker2d	HalfCheetah	Hopper	Walker2d	HalfCheetah
PPO	1292 \pm 650	2196 \pm 200	1127 \pm 308	0.5	0.5	0.5
WSR	1253 \pm 591	1992 \pm 380	1091 \pm 469	0.6	0.3	0.3
TNB	1699 \pm 573	1788 \pm 214	887 \pm 178	0.8	0.0	0.1
CPO	1681 \pm 696	2082 \pm 660	1194 \pm 215	0.8	0.6	0.8
CTNB	1721 \pm 765	2405 \pm 177	1251 \pm 473	0.8	0.9	0.5
IPD (Ours)	2536 \pm 557	2282 \pm 206	1875 \pm 533	1.0	0.6	0.9

all better than their multi-objective counterparts from TNB, showing the superiority of generating novel policies with constrained optimization. In all experiments we use a linear novelty reward function, i.e., $r_{\text{int}} = \min_{\theta_j} \bar{D}_W^{\rho}(\theta, \theta_j)$. We attribute the failure of CPO, TNB and CTNB in Walker and HalfCheetah in finding novel policies to that their convergence behavior is fully controlled by the reward scaling function. Whereas in IPD, there is no novelty-gradient controlled by such a scaling function. The IPD method consistently provides more novelty than CTNB and CPO, while the primal task performances are still guaranteed.

Comparisons of the task-related rewards are carried out in Table 1, where among all the four methods, IPD provides sufficient diversity with minimum loss of performance. Instead of performance decay, we find IPD is able to find better policies in the environment of Hopper and HalfCheetah. Moreover, in the Hopper environment, while the agents trained with PPO tend to fall into the same local minimum. (e.g., they all jump as far as possible and then terminate this episode. On the contrary, PPO with IPD keeps new agents away from falling into the same local minimum, because once an agent has reached some local minimum, agents learned later will try to avoid this region due to the novelty constraints. Such property shows that IPD can enhance the traditional RL schemes to tackle the local exploration challenge (Tessler et al., 2019; Ciosek et al., 2019). A similar feature brings about reward growth in the environment of HalfCheetah.

3.2.2 SUCCESS RATE OF EACH METHOD

In addition to averaged reward, we also use the success rate as another metric to compare the performance of different approaches. Roughly speaking, the success rate evaluates the stability of each method in terms of generating a policy that performs as good as the policies PPO generates. In this work, we regard a policy successful when its performance achieves at least as good as the median performance of policies trained with PPO. To be specific, we use the median of the final performance of PPO as the baseline, and if a novel policy, which aims at performing differently to solve the same task, surpasses the baseline during its training process, it will be regarded as a successful policy. By definition, the success rate of PPO is 0.5 as a baseline for every environment. Table 1 shows the success rate of all the methods. The results show that all constrained novel policy generation methods (CTNB, IPD, CPO) can surpass the average baseline during training, while the multi-objective optimization approaches normally can not. Thus the performance of constrained novel policy generation methods can always be insured.

4 CONCLUSION

In this work, we rethink the novel policy seeking problem under the perspective of constrained optimization. We first introduce a new metric to measure the distances between policies, on top of which we define the novelty of a policy. Based on our formulation of constrained optimization, we provide practical algorithms for constrained novel policy learning, we evaluate several constrained policy optimization methods: namely the CPO, Constrained TNB, and the Interior Policy Differentiation (IPD) proposed in this work. Our experimental results demonstrate IPD, as a novelty (policy) gradient-free approach, can effectively learn various well-performing yet diverse policies, outperforming previous multi-objective methods, as well as constrained optimization baselines.

ETHICS STATEMENT

In this work, we study the problem of diverse policy generation. Although in this work we empirically demonstrate the effectiveness of the proposed algorithm and learning framework on simulated reinforcement learning benchmarks, we believe there are lots of real-world applications of reinforcement learning that can be potentially benefited from our research. To name some examples: first and foremost, behavior diversity itself is precious in that it enables reinforcement learning agents to have their characteristics, being able to generate various well-performing agents will improve the interactive experience in games with both human and AI players. Moreover, as learning novel policies can help to escape from local-optimal solutions, it can potentially benefit downstream tasks such as cooperative exploration. Human society indeed needs individual diversity in cooperatively exploring and exploiting knowledge. Last but not least, diversity helps species in evolution from the perspective of *survival of the fittest*, therefore agents trained to finish the task with distinct behavior can potentially benefit the robustness of the population under environmental changes.

REPRODUCIBILITY STATEMENT

We include our code in the supplementary materials. More implementation details of our experiments are provided in Appendix E.

REFERENCES

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 22–31. JMLR. org, 2017.
- Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 214–223, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL <http://proceedings.mlr.press/v70/arjovsky17a.html>.
- Adrià Puigdomènech Badia, Bilal Piot, Steven Kapturowski, Pablo Sprechmann, Alex Vitvitskyi, Zhaohan Daniel Guo, and Charles Blundell. Agent57: Outperforming the atari human benchmark. In *International Conference on Machine Learning*, pp. 507–517. PMLR, 2020.
- Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.
- Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A Efros. Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*, 2018a.
- Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018b.
- Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A lyapunov-based approach to safe reinforcement learning. In *Advances in neural information processing systems*, pp. 8092–8101, 2018.
- Kamil Ciosek, Quan Vuong, Robert Loftin, and Katja Hofmann. Better exploration with optimistic actor critic. In *Advances in Neural Information Processing Systems*, pp. 1785–1796, 2019.

- A Conn, Nick Gould, and Ph Toint. A globally convergent lagrangian barrier algorithm for optimization with general inequality constraints and simple bounds. *Mathematics of Computation of the American Mathematical Society*, 66(217):261–288, 1997.
- Edoardo Conti, Vashisht Madhavan, Felipe Petroski Such, Joel Lehman, Kenneth Stanley, and Jeff Clune. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. In *Advances in Neural Information Processing Systems*, pp. 5027–5038, 2018.
- George B Dantzig and Mukund N Thapa. *Linear programming 2: theory and extensions*. Springer Science & Business Media, 2006.
- Mahmoud Elbarbari, Kyriakos Efthymiadis, Bram Vanderborght, and Ann Nowé. Ltlf-based reward shaping for reinforcement learning. In *Adaptive and Learning Agents Workshop 2021*, 2021.
- Dominik Maria Endres and Johannes E Schindelin. A new metric for probability distributions. *IEEE Transactions on Information theory*, 2003.
- Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- Bent Fuglede and Flemming Topsøe. Jensen-shannon divergence and hilbert space embedding. In *International Symposium on Information Theory, 2004. ISIT 2004. Proceedings.*, pp. 31. IEEE, 2004.
- Scott Fujimoto, Herke Van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*, 2018.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- Yuval Noah Harari. *Sapiens: A brief history of humankind*. Random House, 2014.
- Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Joseph Henrich. *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press, 2017.
- José Herskovits. Feasible direction interior-point technique for nonlinear optimization. *Journal of optimization theory and applications*, 99(1):121–146, 1998.
- Zhang-Wei Hong, Tzu-Yun Shann, Shih-Yang Su, Yi-Hsiang Chang, Tsu-Jui Fu, and Chun-Yi Lee. Diversity-driven exploration strategy for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 10489–10500, 2018.
- Rein Houthooft, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. Variational information maximizing exploration. 2016.
- Edward Hughes, Joel Z Leibo, Matthew Phillips, Karl Tuyls, Edgar Dueñez-Guzman, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Koster, et al. Inequity aversion improves cooperation in intertemporal social dilemmas. In *Advances in neural information processing systems*, pp. 3326–3336, 2018.
- Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro Ortega, Dj Strouse, Joel Z Leibo, and Nando De Freitas. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *International Conference on Machine Learning*, pp. 3040–3049, 2019.
- Adam Daniel Laud. *Theory and application of reward shaping in reinforcement learning*. University of Illinois at Urbana-Champaign, 2004.

- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Hao Liu, Alexander Trott, Richard Socher, and Caiming Xiong. Competitive experience replay. *CoRR*, abs/1902.00528, 2019. URL <http://arxiv.org/abs/1902.00528>.
- Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 16–17, 2017.
- Alexander Peysakhovich and Adam Lerer. Consequentialist conditional cooperation in social dilemmas with imperfect information. *arXiv preprint arXiv:1710.06975*, 2017.
- Florian A Potra and Stephen J Wright. Interior-point methods. *Journal of Computational and Applied Mathematics*, 124(1-2):281–302, 2000.
- Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.
- Jette Randsløv and Preben Alstrøm. Learning to drive a bicycle using reinforcement learning and shaping. In *ICML*, volume 98, pp. 463–471. Citeseer, 1998.
- Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. *openai*, 2019.
- Barbara Rogoff. *Apprenticeship in thinking: Cognitive development in social context*. Oxford university press, 1990.
- Ludger Rüschendorf. The wasserstein distance and approximation theorems. *Probability Theory and Related Fields*, 70(1):117–129, 1985.
- Andrzej Ruszczyński. Feasible direction methods for stochastic programming problems. *Mathematical Programming*, 19(1):220–229, 1980.
- Richard M Ryan and Edward L Deci. Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary educational psychology*, 25(1):54–67, 2000.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Pedro Sequeira, Francisco S Melo, Rui Prada, and Ana Paiva. Emerging social awareness: Exploring intrinsic motivation in multiagent learning. In *2011 IEEE International Conference on Development and Learning (ICDL)*, volume 2, pp. 1–6. IEEE, 2011.
- Archit Sharma, Shixiang Gu, Sergey Levine, Vikash Kumar, and Karol Hausman. Dynamics-aware unsupervised discovery of skills. *arXiv preprint arXiv:1907.01657*, 2019.
- Felipe Petroski Such, Vashisht Madhavan, Rosanne Liu, Rui Wang, Pablo Samuel Castro, Yulun Li, Ludwig Schubert, Marc Bellemare, Jeff Clune, and Joel Lehman. An atari model zoo for analyzing, visualizing, and comparing deep reinforcement learning agents. *arXiv preprint arXiv:1812.07069*, 2018.
- Hao Sun, Ziping Xu, Meng Fang, Zhenghao Peng, Jiadong Guo, Bo Dai, and Bolei Zhou. Safe exploration by solving early terminated mdp. *arXiv preprint arXiv:2107.04200*, 2021.
- Richard S Sutton, Andrew G Barto, et al. *Introduction to reinforcement learning*, volume 2. MIT press Cambridge, 1998.
- Chen Tessler, Guy Tennenholtz, and Shie Mannor. Distributional policy optimization: An alternative approach for continuous control. *arXiv preprint arXiv:1905.09855*, 2019.

- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *IROS*, pp. 5026–5033. IEEE, 2012. ISBN 978-1-4673-1737-5. URL <http://dblp.uni-trier.de/db/conf/iros/iros2012.html#TodorovET12>.
- Carel P van Schaik and Judith M Burkart. Social learning and evolution: the cultural intelligence hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1567): 1008–1016, 2011.
- Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Stephen J Wright. On the convergence of the newton/log-barrier method. *Mathematical Programming*, 90(1):71–100, 2001.
- Yunbo Zhang, Wenhao Yu, and Greg Turk. Learning novel policies for tasks. In *International Conference on Machine Learning*, pp. 7483–7492. PMLR, 2019.

APPENDIX

A METRIC SPACE

Definition 2. A metric space is an ordered pair (M, d) where M is a set and d is a metric on M , i.e., a function $d: M \times M \rightarrow \mathbb{R}$ such that for any $x, y, z \in M$, the following holds:

1. $d(x, y) \geq 0, d(x, y) = 0 \Leftrightarrow x = y$,
2. $d(x, y) = d(y, x)$,
3. $d(x, z) \leq d(x, y) + d(y, z)$.

B PROOF OF PROPOSITION 1

The first two properties are obviously guaranteed by \overline{D}_W^ρ . As for the triangle inequality,

$$\begin{aligned}
& \mathbb{E}_{s \sim \rho(s)} [D_W(\theta_i(s), \theta_k(s))] \\
&= \mathbb{E}_{s \sim \rho(s)} \left[\sum_{l=1}^{|\mathcal{A}|} |\theta_i(s) - \theta_k(s)| \right] \\
&= \mathbb{E}_{s \sim \rho(s)} \left[\sum_{l=1}^{|\mathcal{A}|} |\theta_i(s) - \theta_j(s) + \theta_j(s) - \theta_k(s)| \right] \\
&\leq \mathbb{E}_{s \sim \rho(s)} \left[\sum_{l=1}^{|\mathcal{A}|} (|\theta_i(s) - \theta_j(s)| + |\theta_j(s) - \theta_k(s)|) \right] \\
&= \mathbb{E}_{s \sim \rho(s)} \left[\sum_{l=1}^{|\mathcal{A}|} |\theta_i(s) - \theta_j(s)| \right] + \mathbb{E}_{s \sim \rho(s)} \left[\sum_{l=1}^{|\mathcal{A}|} |\theta_j(s) - \theta_k(s)| \right] \\
&= \mathbb{E}_{s \sim \rho(s)} [D_W(\theta_i(s), \theta_j(s))] + \mathbb{E}_{s \sim \rho(s)} [D_W(\theta_j(s), \theta_k(s))]
\end{aligned}$$

C PROOF OF PROPOSITION 2

$$\begin{aligned}
\rho_\theta(s) &= P(s_0 = s|\theta) + P(s_1 = s|\theta) + \dots + P(s_T = s|\theta) \\
&\stackrel{L.L.N.}{=} \lim_{N \rightarrow \infty} \frac{\sum_{i=1}^N I(s_0 = s|\tau_i)}{N} + \frac{\sum_{i=1}^N I(s_1 = s|\tau_i)}{N} + \dots + \frac{\sum_{i=1}^N I(s_T = s|\tau_i)}{N} \\
&= \lim_{N \rightarrow \infty} \frac{\sum_{j=0}^T \sum_{i=1}^N I(s_j = s|\tau_i)}{N} \\
\bar{\rho}_\theta(s) &= \sum_{i=1}^N \sum_{j=0}^T \frac{I(s_j = s|\tau_i)}{N} \\
\mathbb{E}[\bar{\rho}_\theta(s) - \rho_\theta(s)] &= 0
\end{aligned}$$

D MORE DETAILS ON WSR

WSR: Penalty Method The Penalty Method considers the constraints of Eq.(6) by putting constraint $g(\theta)$ into a penalty term, followed by solving the following unconstrained problem in an iterative manner,

$$\max_{\theta \in \Theta} f(\theta) + \frac{1-\alpha}{\alpha} \min\{g(\theta), 0\}, \quad (8)$$

The limit of the above unconstrained problem when $\alpha \rightarrow 0$ then leads to the solution of the original constrained problem. As an approximation, WSR chooses a fixed weight α , and uses the gradient of $\nabla_\theta f + \frac{1-\alpha}{\alpha} \nabla_\theta g$ instead of $\nabla_\theta f + \frac{1-\alpha}{\alpha} \nabla_\theta \min\{g(\theta), 0\}$, thus the final solution will intensely rely on the selection of α .




Figure 4: The performance under different novelty thresholds in the Hopper, Walker and HalfCheetah environments. The results are collected from 10 learned policies based on PPO. The box extends from the lower to upper quartile values of the data, with a line at the median. The whiskers extend from the box to show the range of the data. Flier points are those past the end of the whiskers.

E IMPLEMENTATION DETAILS

Calculation of D_W We use deterministic part of policies in the calculation of D_W , i.e., we remove the Gaussian noise on the action space in PPO and use $D_W(a_1, a_2) = |a_1 - a_2|$.

Network Structure We use MLP with 2 hidden layers as our actor models in PPO. The first hidden layer is fixed to have 32 units. We choose to use 10, 64 and 256 hidden units for the three tasks respectively in all of the main experiments, after taking the success rate, performance and computation expense (i.e. the preference to use less unit when the other two factors are similar) into consideration.

Training Timesteps We fix the training timesteps in our experiments. The timesteps are fixed to be 1M in Hopper-v3, 1.6M for Walker2d-v3 and 3M for HalfCheetah-v3.

Hardware and Training Time We experiment on a server with 8 TITAN X GPUs and 32 Intel(R) E5-2640 CPUs. Experiments take 0.5 (the maze environment with 0.1M interactions) to 10 hours (for the locomotion tasks) to run in generating each policy. As our diversity-seeking method based on a sequential learning framework, new policies must be trained after previous policies are already generated.

F ABLATION STUDY

Novel Policy Generation without Performance Decay Multi-objective formulation of novel policy generation has the risk of sacrificing the primal performance as the overall objective needs to consider both novelty and primal task rewards. On the contrary, under the perspective of constrained optimization, there will be no more trade-off between novelty and final reward as the only objective is the task reward. Given a certain novelty threshold, the algorithms tend to find the optimal solution in terms of task reward under constraints, thus the learning process becomes more controllable and reliable, i.e., one can utilize the novelty threshold to control the degree of novelty.

The proper magnitude of the novelty threshold leads to more exploration among a population of policies, thus the performance of latterly found policies may be better than or at least as good as those trained without novelty seeking. However, when a larger magnitude of novelty threshold is applied, the performance of found novel policies will decrease because finding a feasible solution will get harder under more strict constraints. Fig. 4 shows experimental results on adjusting the thresholds, which supports our intuition.

G VISUALIZE DIVERSITY

G.1 MUJOCO LOCOMOTION

In this section, we provide some qualitative results of IPD on the Mujoco locomotion tasks. In all of our experiments we use the vanilla Mujoco locomotion benchmarks, with the default settings on defining healthy states. Although otherwise the visualization of learned policies might become more

diverse (e.g., a Hopper agent may learn to stand-up after falling down while another agent may learn to move forward on the ground if we set the z -axis healthy threshold as 0).

With the method of IPD, the Hopper policies (Figure 5) learns to jump further and avoids falling down rather instead of just jumping and falling down (Figure 6). In the Walker2d environment, the color of purple indicates the left leg is visible. It can be seen that the IPD policies (Figure 7) learn to use both left and right legs in walking, while the PPO policies usually learn jumping. (Figure 8). In HalfCheetah, the IPD policies (Figure 9) perform much better than the PPO policies (Figure 10). The IPD policies learn to run with head-downward (Figure 9 line 1), head-upward (Figure 9 line 3), and forward (Figure 9 line 5) while the PPO policies are always head-downward.

In Hopper and HalfCheetah, IPD is able to improve the primal task performance by avoiding always getting trapped in some certain sub-optimal behaviors.




Figure 5: The visualization of policy behaviors of agents trained by our method in Hopper-v3 environment. Agents learn to jump with different strides.




Figure 6: The visualization of policy behaviors of agents trained by PPO in Hopper-v3 environment. Most agents learn a policy that can be described as *Jump as far as possible and fall down*, leading to relative poor performance.



Figure 7: The visualization of policy behaviors of agents trained by our method in Walker2d-v3 environment. Instead of bouncing at the ground using both legs, our agents learn to use both legs to step forward.




Figure 8: The visualization of policy behaviors of agents trained by PPO in Walker2d-v3 environment. Most of the PPO agents only learn to use their right leg to support the body and jump forward.




Figure 9: The visualization of policy behaviors of agents trained by our method in HalfCheetah-v3 environment. Our agents run much faster compared to PPO agents and at the mean time several patterns of motion have emerged.




Figure 10: The visualization of policy behaviors of agents trained by PPO in HalfCheetah-v3 environment. Since we only draw fixed number of frames in each line, in the limited time steps the PPO agents can not run enough distance to leave the range of our drawing, which shows that our agents run much faster.