

Diversity-Inducing Policy Gradient: Using Maximum Mean Discrepancy to Find a Set of Diverse Policies

Muhammad Masood* and Finale Doshi-Velez

Harvard University

masood@g.harvard.edu, finale@seas.harvard.edu

Abstract

Standard reinforcement learning methods aim to master one way of solving a task whereas there may exist multiple near-optimal policies. Being able to identify this *collection* of near-optimal policies can allow a domain expert to efficiently explore the space of reasonable solutions. Unfortunately, existing approaches that quantify uncertainty over policies are not ultimately relevant to finding policies with qualitatively distinct behaviors. In this work, we formalize the difference between policies as a difference between the distribution of trajectories induced by each policy, which encourages diversity with respect to both state visitation and action choices. We derive a gradient-based optimization technique that can be combined with existing policy gradient methods to now identify diverse collections of well-performing policies. We demonstrate our approach on benchmarks and a healthcare task.

1 Introduction

Standard reinforcement learning methods find one way to solve a task, even though there may exist multiple near-optimal policies that are distinct in some meaningful way. Identifying this *collection* of near-optimal policies can allow a domain expert to efficiently explore the space of reasonable solutions. For example, knowing that there exist comparably-performing policies that trade between several small doses of a drug or a single large dose may enable a clinician to identify what might work best for a patient (e.g., based on whether the patient will remember to take all the small doses).

Unfortunately, existing approaches to find a set of diverse policies involve notions of diversity that are not aligned with the kind of efficient exploration-amongst-reasonable-options setting described above. [Liu *et al.*, 2017] characterize the uncertainty over policies via computing a posterior over policy parameters, but differences in policy parameters may not result in qualitatively different behavior (especially in over-parameterized architectures). [Haarnoja *et al.*, 2017] encourage diversity via encouraging high entropy distributions over

actions (given states), which may result in sub-optimal behavior. [Fard and Pineau, 2011] seek a single non-deterministic policy that may make multiple decisions at any state, which may be overly restrictive if action choices across states must be correlated to achieve near-optimal performance.

We argue that differences in trajectories (state visits and action choices) better capture the kinds of distinct behavior we are seeking. For example, does one prefer a policy that achieves wellness through a surgery, or via prolonged therapy? More formally, stochasticity in the environment dynamics and the policy will induce a distribution over trajectories. We use the maximum mean discrepancy (MMD) metric to compare these distributions over trajectories under different policies. As noted in [Sriperumbudur *et al.*, 2010], the MMD metric has a closed form solution (unlike the Wasserstein and Dudley metrics) and exhibits better convergence behavior than ϕ -divergences such as Kullback-Leibler (KL).

In this work, we first formalize notions of policy diversity via the MMD over their induced trajectories. We show that unbiased gradient estimates of the MMD term can be obtained without knowledge of transition dynamics of the environment, and describe how it can be applied to *any* policy gradient objective. Across both benchmark domains and a healthcare task, our approach discovers diverse collections of well-performing policies.

2 Background

Reinforcement Learning

We consider Markov decision processes (MDPs) defined by a continuous state space \mathcal{S} , a (discrete or continuous) action space \mathcal{A} , state transition probabilities $p_T(s, a, s')$, reward function $r(s, a)$, as well as a discount factor γ . A policy $\pi(s, a)$ indicates the probability of action a in state s ; together with the transition probability $p_T(s, a, s')$, it induces a distribution $p_\pi(\tau)$ over trajectories $\tau = s_0, a_0, \dots, a_{T-1}, s_T$. Traditionally, the task is to find a policy that maximizes the long-term expected discounted sum of rewards (return) $g(\tau)$:

$$g(\tau) = \mathbb{E}_{\tau \sim p_\pi(\tau)} \left[\sum_{t=0}^T \gamma^t r(s_t, a_t) \right]$$

In this work, we shall seek multiple near-optimal policies.

*Contact Author

Policy Gradient Methods

Policy gradient methods use gradients to iteratively optimize a policy $\pi_\theta(s, a)$ that is parameterized by θ . The standard objective $J_{PG}(\theta)$ is the return $g(\tau)$. An unbiased estimate of the gradient of the objective can be obtained by Monte Carlo rollouts generated by the policy π_θ using the likelihood ratio trick. For a single rollout $\{s_t, a_t, r_t\}$, the gradient can be estimated as

$$\nabla_\theta J_{PG}(\theta) \approx \sum_{t=0}^T \nabla_\theta \log \pi_\theta(s_t, a_t) g_t$$

where $g_t = \sum_{t'=t}^T \gamma^{t'-t} r_{t'}$ is the return over the rewards received from time t onwards. [Schulman *et al.*, 2017] introduced Proximal Policy Optimization (PPO); a popular variant that achieves state-of-the-art performance on many benchmark domains and uses trust region updates that are compatible with stochastic gradient descent.

In this work we augment the PPO objective for the policy network with an MMD-based term which encourages policies that lead to different state visitation and/or action choice distribution than previously-identified policies. An iterative use of this diversity inducing variant of policy gradient methods allows us to sequentially obtain distinct policies.

Off-Policy Evaluation

The off-policy RL framework applies when trajectories are collected from a behavior policy that is distinct from the policy being trained, e.g., using observations from clinician behavior to optimize an agent’s behavior. In the batch setting, there are model-based and importance sampling approaches to estimating the value of a policy. Model-based approaches learn a dynamics model for the transitions and use those to simulate outcomes of a policy in order to estimate its value. Importance sampling approaches appropriately re-weight existing batch data to estimate the value for a different policy.

Maximum Mean Discrepancy

We use the MMD metric to measure the difference between two trajectory distributions. The MMD is an integral probability metric [Gretton *et al.*, 2007] that measures the difference between two distributions p, q using test functions h from a function space \mathcal{H} . It is given by:

$$\text{MMD}(p, q, \mathcal{H}) = \sup_{h \in \mathcal{H}} (\mathbb{E}_{x \sim p}[h(x)] - \mathbb{E}_{y \sim q}[h(y)])$$

Computing the MMD is tractable when the function space \mathcal{H} is a unit-ball in a reproducing kernel hilbert space (RKHS) defined by a kernel $k(\cdot, \cdot)$ and is given by:

$$\text{MMD}^2(p, q, \mathcal{H}) = \mathbb{E}[k(x, x')] - 2\mathbb{E}[k(x, y)] + \mathbb{E}[k(y, y')]$$

where x, x' i.i.d. $\sim p$ and y, y' i.i.d. $\sim q$. The expectation terms in the analytical expression for the MMD can be approximated using samples. In Section 3, we will show that computing the derivative of the MMD metric between trajectories with respect to the policy parameters θ is tractable.

3 Diversity-Inducing Policy Gradient (DIPG)

Our algorithm constructs a set of diverse policies for an MDP by iteratively finding policies that are diverse relative to an existing set of policies. First, we formulate a diversity inducing

objective function that regularizes any policy gradient objective. Then, we show that optimizing this objective is tractable using the familiar log-derivative trick. Next, we explain how to iteratively apply this diversity-inducing policy gradient objective in conjunction with an existing algorithm to find a set of distinct policies that solve an MDP. Finally, we introduce an extension to the proposed framework for off-policy batch reinforcement learning.

3.1 DIPG Objective

We propose adding a regularization term to encourage learning a policy $\pi_\theta(s, a)$ that induces a distribution over trajectories $p_\theta(\tau)$ that is distinct from a specified set of distributions over trajectories $\mathcal{Q} = \{q_m(\tau)\}_{m=1}^M$. (Below, these distributions $q_m(\tau)$ will come from previously-identified policies.) Our diversity measure $D_{\text{MMD}}(p_\theta(\tau), \mathcal{Q})$ is the squared MMD between the distribution of trajectories under the current policy π_θ and the distribution $q_*(\tau)$ in \mathcal{Q} most similar to it.

$$\begin{aligned} D_{\text{MMD}}(p_\theta(\tau), \mathcal{Q}) &= \min_{m \in \{1, \dots, M\}} \text{MMD}^2(p_\theta(\tau), q_m(\tau)) \\ &= \text{MMD}^2(p_\theta(\tau), q_*(\tau)) \end{aligned}$$

We define a kernel over a pair of trajectories (τ, τ') using a kernel K (such as Gaussian kernel) defined over some function (g) of a vectorized representation of N steps of trajectories ($x = \phi_N(\tau), x' = \phi_N(\tau')$). For trajectories that are not of the same length, we pick N to be the number of steps corresponding to the shorter of the two trajectories.

$$k(\tau, \tau') = K(g(\phi_N(\tau)), g(\phi_N(\tau')))$$

The function ϕ_N is simply a way to stack the states and actions from the first N steps of a trajectory into a single vector. The function g gives the flexibility to adjust the focus of the diversity measure for different aspects such as state visits, action choices, or both.

The MMD-based diversity-inducing objective function is given by $J_{\text{MMD}}(\theta)$:

$$J_{\text{MMD}}(\theta) = J_{\text{PG}}(\theta) + \alpha D_{\text{MMD}}(p_\theta(\tau), \mathcal{Q})$$

Where $J_{\text{PG}}(\theta)$ is the objective function of a policy gradient algorithm of the user’s choice (e.g. for vanilla policy gradient it would be the expected return) and the second term is proportional to the diversity measure between a policy distribution and the set of specified distributions \mathcal{Q} . The parameter α decides the relative importance of optimality and diversity of the new policy $\pi_\theta(s, a)$.

To optimize the MMD-based diversity-inducing objective, we need to specify how gradients of the diversity term can be computed with respect to the policy parameters θ .

3.2 Optimization via Gradient Ascent

To use gradient ascent on $J_{\text{MMD}}(\theta)$, what remains to specify is the gradient with respect to θ of the diversity inducing term D_{MMD} . Let $q_*(\tau)$ be the distribution in \mathcal{Q} that minimizes the MMD between the state trajectory distribution of the policy π_θ and $q_m \in \mathcal{Q}$. Then, the gradient with respect to the policy

Algorithm 1 MMD-based Diversity-Inducing Policy

Input: Known policies $\mathcal{P}_{\text{known}}$, MDP $\{\mathcal{S}, \mathcal{A}, p_s, r\}$, policy gradient objective J_{PG} , learning rate η , diversity parameter α
 Initialize policy parameters θ
 $\mathcal{Q} \leftarrow$ Sampled trajectories from policies in $\mathcal{P}_{\text{known}}$
repeat
 Generate an episode $s_0, a_0, r_1, \dots, s_{T-1}, a_{T-1}, r_T$, following policy π_θ
 for Each step $t = 0, \dots, T - 1$ **do**
 1: Estimate $\nabla_\theta J_{\text{PG}}$ and $\nabla_\theta D_{\text{MMD}}(p_\theta(\tau), \mathcal{Q})$
 2: Update policy parameters via gradient ascent
 $\theta \leftarrow \theta + \eta(\nabla_\theta J_{\text{PG}} + \alpha \nabla_\theta D_{\text{MMD}}(p_\theta(\tau), \mathcal{Q}))$
 end for
until convergence
Output: policy p_θ

Algorithm 2 DIPG

Input: Number of policies N , MDP $\{\mathcal{S}, \mathcal{A}, p_s, r\}$, policy gradient objective J_{PG} , learning rate η , diversity parameters $\alpha_{1 \dots N}$
 Collection of known policies $\mathcal{P}_{\text{known}} = \emptyset$
for $n = 1$ **to** N **do**
 1: Find a policy p_n that is distinct from the current set of known policies $\mathcal{P}_{\text{known}}$:
 $p_n \leftarrow$ Algorithm 1($\mathcal{P}_{\text{known}}, \text{MDP}, J_{\text{PG}}, \eta, \alpha_n$)
 2: Add p_n to the set of known policies $\mathcal{P}_{\text{known}}$:
 $\mathcal{P}_{\text{known}} \leftarrow \mathcal{P}_{\text{known}} \cup p_n$
end for
Output: Set of policies $\mathcal{P}_{\text{known}}$

parameters θ of the diversity term is given by

$$\begin{aligned} \nabla_\theta D_{\text{MMD}}(p_\theta(\tau), \mathcal{Q}) &= \nabla_\theta \text{MMD}^2(p_\theta(\tau), q_*(\tau), \mathcal{H}) \\ &= \mathbb{E}[k(\tau_p, \tau'_p) \nabla_\theta \log(p_\theta(\tau_p) p_\theta(\tau'_p))] \\ &\quad - 2\mathbb{E}[k(\tau_p, \tau_q) \nabla_\theta \log(p_\theta(\tau_p) q_*(\tau_q))] \\ &\quad + \mathbb{E}[k(\tau_q, \tau'_q) \nabla_\theta \log(q_*(\tau_q) q_*(\tau'_q))] \end{aligned}$$

where τ_p, τ'_p i.i.d. $\sim p_\theta(\tau)$ and τ_q, τ'_q i.i.d. $\sim q_*(\tau)$. The last term only involves the distribution $q_*(\tau) \in \mathcal{Q}$ that has no dependence on θ . The gradient term can be estimated by linear combinations of the $\nabla_\theta \log p_\theta(\tau)$ involving the kernel between sample trajectories from the policy π_θ and with a set of specified trajectories from \mathcal{Q} . It is well-known [Sutton *et al.*, 2000] that the gradient of the score function of the trajectory distribution $\nabla_\theta \log p_\theta(\tau_p)$ does not require the dynamics model and can be expressed in terms of the score function of the policy network ($\nabla_\theta \log p_\theta(\tau) = \sum_{t=0}^H \nabla_\theta \log \pi_\theta(a_t | s_t)$).

We now have all the machinery in place to augment any existing policy gradient method with a diversity inducing term. Next, we will specify the basic algorithm for finding a policy that is diverse with respect to some specified set of trajectory distributions and then introduce an algorithm that leverages this to find a set of diverse policies.

3.3 Finding Multiple Diverse Policies

Our goal is to find a collection of policies that perform optimally or near-optimally and are diverse in terms of the distributions over trajectories that they induce. In Algorithm 1, we state how the diversity inducing term D_{MMD} is used in order to learn a single policy that is distinct with respect to a specified set of distributions over trajectories \mathcal{Q} . In Algorithm 2, we iteratively apply Algorithm 1 to find the desired set of N distinct policies (agents): The first policy is learned without any diversity term because there is no set of known policies to begin with. Subsequent policies are learned by initializing random policy parameters and training with the augmented objective function that encourages diversity with respect to previously discovered policies. The strength of the diversity parameter α can be varied (e.g., to seek diverse policies more aggressively as more policies are discovered).

3.4 DIPG Extension: Batch Off-Policy Setting

Batch reinforcement learning [Lange *et al.*, 2012] aims to learn policies from a fixed set of previously-collected trajectories. It is common in domains such as medicine, dialogue management, and industrial plant control where logged data are plentiful but exploration is expensive or infeasible.

We now extend our DIPG framework to this batch setting. In the on-policy case, we defined the DIPG diversity term as a kernel over the distribution of trajectories induced by different policies. However, in the batch case, trajectories cannot be simulated from the learned policies. Thus, we instead define the diversity as a kernel over the likelihoods of specific (observed) trajectories in the batch with respect to a policy.

Specifically, let $\mathcal{T} = \{\tau_i\}_{i=1}^I$ be a batch of I trajectories. We use $p(\mathcal{T}|\pi) \in \mathbb{R}^I$ to indicate a vector where the i^{th} coordinate equals the probability of the i^{th} trajectory under the policy π i.e. $p_i(\mathcal{T}|\pi) = p(\tau_i|\pi)$. Now, the diversity term can be defined in an analogous fashion to the on-policy case where we compare our policy being optimized π_θ to the previous policies $\{q_1, \dots, q_M\}$:

$$\begin{aligned} D_{\text{BATCH}}^{\mathcal{T}}(\pi_\theta, \mathcal{Q}) &= \min_{m \in \{1, \dots, M\}} k(p(\mathcal{T}|\pi_\theta), p(\mathcal{T}|q_m)) \\ &= k(p(\mathcal{T}|\pi_\theta), p(\mathcal{T}|q_m^*)) \end{aligned}$$

We also require a measure of quality for the policies. [Levine and Koltun, 2013] note that gradient-based optimization of importance sampling estimates is difficult with complex policies and long rollouts. We suggest a surrogate that is equal to the sum of the likelihoods of each trajectory in the batch $J_{\text{Surrogate}}^{\mathcal{T}}(\theta) = \sum_{i=1}^I p(\tau_i|\theta)$. This surrogate is more robust to optimize and encourages ‘safe’ behavior from the agent, a desirable feature in the healthcare setting.

$$J_{\text{BATCH}}(\theta) = J_{\text{Surrogate}}^{\mathcal{T}}(\theta) + \alpha D_{\text{BATCH}}^{\mathcal{T}}(\pi^\theta, \mathcal{Q})$$

While we optimize with respect to this surrogate, in our results, we still report on the value of the policy with respect to a standard importance sampling-based estimator (CWPDIS, from [Thomas, 2015]).

4 Experimental Setup

We augment Proximal Policy Optimization (PPO) [Schulman *et al.*, 2017] with the diversity inducing term (referred to as DIPG-PPO). We set the maximum number of steps taken by all baseline algorithms to be 1 million and set the maximum number of steps for each of the N policies in DIPG-PPO to 0.2 million. Since we choose $N < 5$ for all our testing environments, the total number of steps to learn all N policies is necessarily less than 1 million. In these experiments we set $\alpha_n = 1.0$ across all environments and iterations n^1 . For the MMD kernel, we choose the Gaussian kernel with bandwidth set to 1. Each algorithm is run 3 times to assess variability in performance, except for RR-PPO which does not have any inherent diversity component, so we run it with 10 random restarts to help it make up for this disadvantage.

4.1 Environments

Synthetic Environments

We illustrate qualitative aspects of our approach on 2-dimensional navigational tasks.

Multi-goal environment of [Haarnoja *et al.*, 2017] is solved by reaching one of four symmetrically placed goals in a continuous 2-D world. An ideal collection of diverse policies would include policies that solve the task via reaching different goals in different parts of the state space.

Asymmetric goals environment has one goal that is closer to the starting point and easier to reach than the second. While most agents find the region closest to the initial position, a distinct collection of policies solves the environment by also reaching the goal further away. We use this environment to explore the case where there exist a slightly sub-optimal solution that is clearly distinct from the optimal one.

Benchmark Environments

We evaluated the performance of our algorithm on standard benchmark environments: *Reacher*, *Ant*, *Humanoid* and *Humanoid Flag Run* [Schulman *et al.*, 2017]².

Clinical batch

Hypotension batch data is built from a cohort of patients from MIMIC critical care data set [Johnson *et al.*, 2017]. In this work, we aim to learn multiple treatment strategies using off-policy methods. Cohort and MDP formulation details are in [Masood and Doshi-Velez, 2019].

4.2 Algorithm and Baselines

We compare our approach [DIPG-PPO] to the following other algorithms that can encourage diverse behavior:

Random Restarts [RR-PPO]: Independent runs of PPO without the diversity inducing regularization. Due to variability in initializations and experience, we obtain a collection of policies that are not exactly the same.

Stein Variational Policy Gradient [SVPG]: [Liu *et al.*, 2017] use functional gradient descent to compute a point-based posterior distribution over the policy parameter space.

¹The performance of our algorithm may improve when this parameter is tuned based on the environment or which of the N policies is being learned)

²<https://github.com/benelot/pybullet-gym>

We report results for a collection of 8 agents, and for comparison also on a single agent’s performance.

Deep Energy-Based Policy [Soft-Q]: [Haarnoja *et al.*, 2017] present a way to train a single stochastic policy that is encouraged to have high entropy on the probability of the action (given state). The higher entropy encourages the agent to choose actions that are diverse. We try Soft-Q learning with the default entropy regularization parameter of 1 as well as variants with 0.5 and 0. We choose the smaller (and zero) values of the regularization parameter to see the effect of the tradeoff between diversity and quality in this algorithm.

5 Results

We evaluate the performance based on both quality and similarity scores (figure 2). The quality is evaluated by averaging the returns coming from rollouts of the final stochastic policy (or collection of policies) and the similarity score comes from the kernel used to measure similarity between trajectories (smaller kernel values indicate higher diversity). Additional results can be found in [Masood and Doshi-Velez, 2019].

DIPG method does not sacrifice quality. Figure 2 shows that DIPG-PPO and RR-PPO generally provide the highest quality policies. The baseline algorithms designed for finding a diverse set of policies (SVPG and Soft-Q) have significantly worse quality performance (SVPG in particular) and find policies that lead to high variability in the rewards even though each individual policy in SVPG and Soft-Q is trained for longer (1 million steps) as compared to RR-PPO and DIPG-PPO (0.2 million steps). Such variability in the distribution of returns suggests that the diversity in SVPG and Soft-Q comes not only at the cost of quality, but importantly (especially in the clinical setting) at the cost of consistency of quality returns. It must be emphasized that diversity in policies is only meaningful if they are beyond some threshold of quality and are able to ‘solve’ a domain consistently.

DIPG method finds meaningfully diverse policies with fewer runs. The diversity information in figure 2 shows that SVPG and Soft-Q algorithms give policies that are quite diverse (comparison of pairwise trajectories shows negligible overlap). Since these collections of stochastic policies fail to provide high rewards in the environments, the diversity in the trajectory distribution that is induced is of little value. The DIPG-PPO ($N = 2$ and $N = 4$) and RR-PPO collections ($N = 10$) are collections of policies that essentially ‘solve’ these environments and discover the set of ways these environments were designed to be solved (see figure 1). In comparison to RR-PPO, DIPG-PPO requires many fewer agents to discover an appropriate collection of diverse agents.

Clinical Example: A collection of DIPG policies reveals different choices of treatment in the ICU. We learned four policies through our off-policy extension of DIPG using real data from the ICU and evaluated their performance using the CWPDIS importance sampling estimator [Thomas, 2015]. Not only were all four learned policies of slightly higher quality (~ 9.33 on held out test data) than the clinician behavior policy (9.29), they are also distinct.

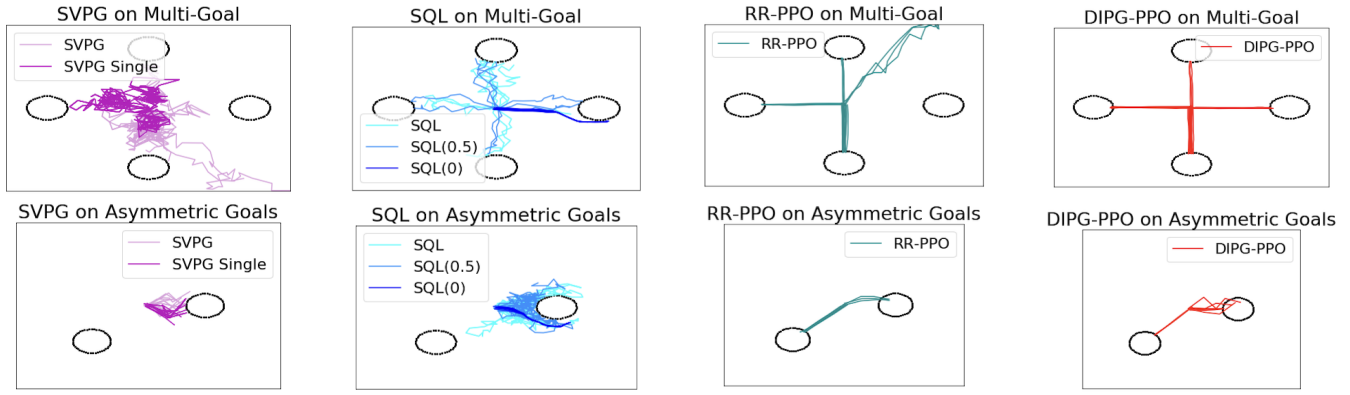


Figure 1: A comparison of trajectories in the 2-D navigation tasks shows that DIPG-PPO (with $N = 4$ for Multi-Goal and $N = 2$ for Asymmetric Goals) produces near-optimal trajectories to reach each goal. Soft-Q takes sub-optimal paths to the goal regions and SVPG generally fails to solve the environments.

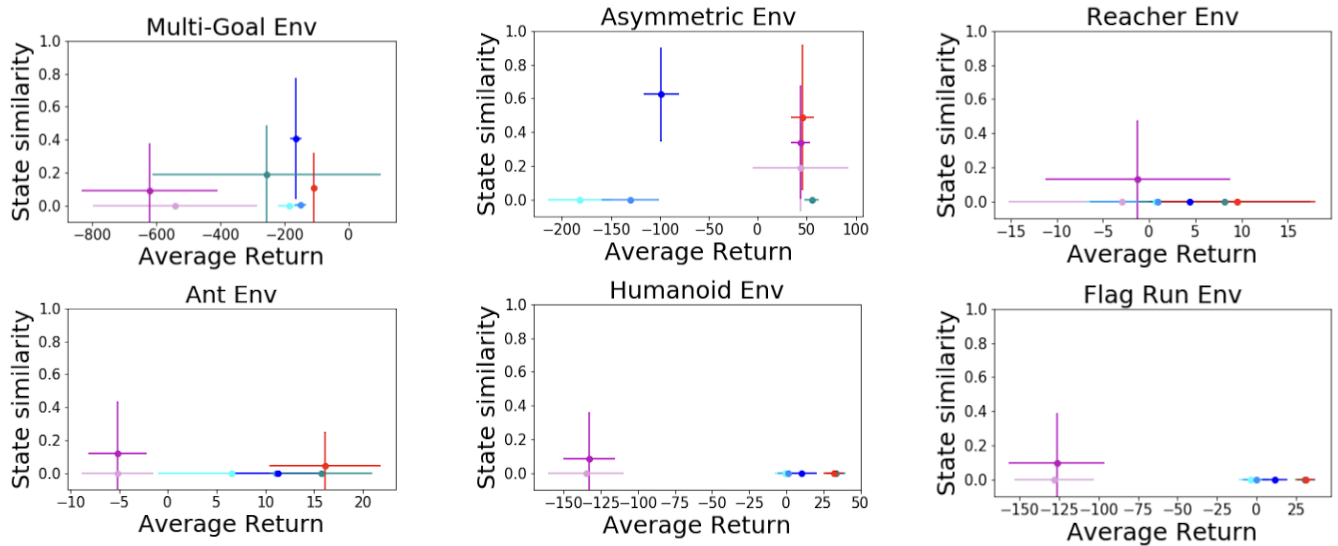


Figure 2: A comparison of the kernel-based similarity between trajectories and the average return across different algorithms reveals that DIPG-PPO and RR-PPO consistently give the largest returns while achieving low similarity. The colors are consistent with those in figure 1

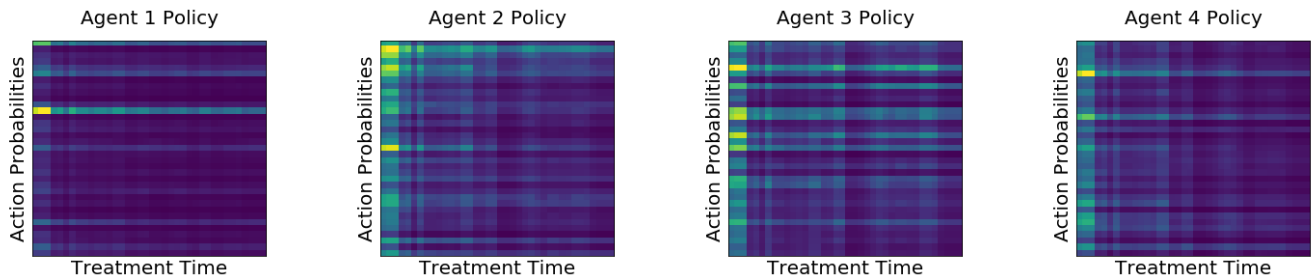


Figure 3: We show the probability of taking actions throughout the duration of a single ICU stay. While these agents all have roughly equal value, they exhibit different emphasis on treatment styles as can be seen by the variability in action probabilities.

Figure 3 plots the probabilities of each action for each policy for a single patient over their ICU stay duration. The discrete actions correspond to different combinations of vasopressor and IV fluid dosages. Agent 1 has a near-deterministic policy with an emphasis on a single treatment that is a combination of a low vasopressor dosage and a medium fluid dosage. Agents 2 and 3 are more aggressive than other agents, giving weights to medium and high vasopressor dosages as well. All agents emphasize taking action early rather than later in the duration (by which time hopefully the patient is in a stable state). Importantly, all of these policies cannot be differentiated by value alone; thus they form a collection to be presented to clinicians for further review of potentially valuable options.

6 Related Work

Whereas we compute diversity (using MMD) between *complete* trajectories (including state and action), many related works, quantify diversity only via differences in the action space. E.g., the KL diversity term in [Hong *et al.*, 2018] is between the actions. Similarly, in [Gangwani *et al.*, 2018], the kernel incorporates trajectory information but the definition of the target density is still according to the action-based maximum entropy RL framework.

[Smith *et al.*, 2018] learn a policy over options and can train multiple options (in an off-policy manner) using a roll-out from a single option. The distinct options can give rise to behavior that is diverse, however there is no explicit diversity component in the objective and it is unclear how to summarize the kinds of distinct trajectories that are possible. Unlike [Eysenbach *et al.*, 2018], our proposed method does not impose a latent structure defining ‘skills’. [Cohen *et al.*, 2019] uses diversity for exploration in order to get closer to an optimal policy whereas our work solves multi-goal domains and identifies diverse policies of interest to experts (e.g. in clinical setting).

In the graph-based planning literature (limited to the discrete setting), there is also an interest in diversity; [Srivastava *et al.*, 2007] seek to do this using domain independent distance measures for evaluating diversity of plans whereas [Sohrabi *et al.*, 2016] first generate a large set of high quality plans and then use clustering to identify a diverse set of representative plans that can be used for further analysis. Our motivation for seeking diverse policies in the reinforcement learning setting is aligned most closely with the end goal of [Sohrabi *et al.*, 2016]: presenting a diverse set of representative solutions as a tool for hypothesis generation and to discover specific directions of interest for further inquiry.

There exists a literature on Bayesian methods for reinforcement learning [Ghavamzadeh *et al.*, 2015]. Also related are evolutionary [Lehman and Stanley, 2011] and multi-objective [Liu *et al.*, 2014] reinforcement learning approaches. However, these approaches do not systematically attempt to identify a small set of distinct policies.

7 Discussion

Our DIPG algorithm successfully finds multiple goals in an optimal or near-optimal manner whereas baseline approaches

are either unable to reach multiple goals or do so sub-optimally. The poor task performance of SVPG could be due to the difficulties of performing functional gradient descent over a high-dimensional parameter space. However, even if those difficulties were to be overcome, the search for diversity in the space of neural network parameters does not correspond directly to any meaningful notion of diversity in the trajectories. The Soft-Q algorithm exhibits some meaningful diversity in the policies it finds (e.g. reaching different goals in the multi-goal environment), however, there is unnecessary stochasticity in the actions, leading to sub-optimal policies. Unlike the baselines, random restarts do result in near-optimal in-task performance. However, the diversity of the policies is at the mercy of the basins of attraction of each local optima (based on the initialization and subsequent experience). In contrast, DIPG-PPO successfully identifies a collection of distinct policies consistently and efficiently (that is, with few runs).

While we have focused on certain notions of distinctness here (state visits, action counts), our approach extends easily to other measures of distinctness as measured by alternative kernels (see [Chen *et al.*, 2018] and [Gretton *et al.*, 2012] for novel kernels and discussion). Whereas the full trajectory may be needed for training the RL agent, we can capture (MMD-based) distinctness more generally using an arbitrary function of the trajectory. For example, we could group patients into clinically meaningful clusters which can be used to define a kernel for measuring diversity. There are also opportunities for incorporating more efficient search algorithms than gradient descent (e.g. [Toussaint and Lopes, 2017]).

While we have focused on the task of returning policies as possible options to a human user, another use-case could be for situations in which we have a cheap, low fidelity simulator and a more expensive, high-fidelity simulator. In this case, the distinct trajectories from the low-fidelity simulator could be used as seeds for deep exploration in the more expensive simulator. Diverse policies can help manage the exploration-exploitation trade-off.

8 Conclusion

We presented an approach for identifying a collection of near-optimal policies with significantly different distributions of trajectories. Our MMD-based regularizer can be applied to the distribution of any statistic of the trajectories—state visits, action counts, state-action combinations—and can also be easily incorporated into any policy-gradient method. We applied these to several standard benchmarks and developed an off-policy extension to identify meaningfully different treatment options from observational clinical data. The ability to find these diverse policies may be useful when the agent does not have complete information about the task, and for presenting a set of potentially reasonable options to a downstream human or system, who can use that information to efficiently choose amongst reasonable options.

Acknowledgements

We acknowledge support from AFOSR FA 9550-17-1-0155.

References

- [Chen *et al.*, 2018] W. Y. Chen, L. Mackey, J. Gorham, F.-X. Briol, and C. J. Oates. Stein Points. *ArXiv e-prints*, March 2018.
- [Cohen *et al.*, 2019] Andrew Cohen, Xingye Qiao, Lei Yu, Elliot Way, and Xiangrong Tong. Diverse exploration via conjugate policies for policy gradient methods. *arXiv preprint arXiv:1902.03633*, 2019.
- [Eysenbach *et al.*, 2018] 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.
- [Fard and Pineau, 2011] Mahdi Milani Fard and Joelle Pineau. Non-deterministic policies in markovian decision processes. *J. Artif. Intell. Res. (JAIR)*, 40:1–24, 2011.
- [Gangwani *et al.*, 2018] Tanmay Gangwani, Qiang Liu, and Jian Peng. Learning self-imitating diverse policies. *arXiv preprint arXiv:1805.10309*, 2018.
- [Ghavamzadeh *et al.*, 2015] Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, Aviv Tamar, et al. Bayesian reinforcement learning: A survey. *Foundations and Trends® in Machine Learning*, 8(5-6):359–483, 2015.
- [Gretton *et al.*, 2007] Arthur Gretton, Karsten M Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex J Smola. A kernel method for the two-sample-problem. In *Advances in neural information processing systems*, pages 513–520, 2007.
- [Gretton *et al.*, 2012] Arthur Gretton, Dino Sejdinovic, Heiko Strathmann, Sivaraman Balakrishnan, Massimiliano Pontil, Kenji Fukumizu, and Bharath K Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. In *Advances in neural information processing systems*, pages 1205–1213, 2012.
- [Haarnoja *et al.*, 2017] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. *arXiv preprint arXiv:1702.08165*, 2017.
- [Hong *et al.*, 2018] 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*, pages 10489–10500, 2018.
- [Johnson *et al.*, 2017] Alistair EW Johnson, David J Stone, Leo A Celi, and Tom J Pollard. The mimic code repository: enabling reproducibility in critical care research. *Journal of the American Medical Informatics Association*, 25(1):32–39, 2017.
- [Lange *et al.*, 2012] Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In *Reinforcement learning*, pages 45–73. Springer, 2012.
- [Lehman and Stanley, 2011] Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, 2011.
- [Levine and Koltun, 2013] Sergey Levine and Vladlen Koltun. Guided policy search. In *International Conference on Machine Learning*, pages 1–9, 2013.
- [Liu *et al.*, 2014] Chunming Liu, Xin Xu, and Dewen Hu. Multiobjective reinforcement learning: A comprehensive overview. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45(3):385–398, 2014.
- [Liu *et al.*, 2017] Yang Liu, Prajit Ramachandran, Qiang Liu, and Jian Peng. Stein variational policy gradient. *arXiv preprint arXiv:1704.02399*, 2017.
- [Masood and Doshi-Velez, 2019] Muhammad A Masood and Finale Doshi-Velez. Diversity-inducing policy gradient: Using maximum mean discrepancy to find a set of diverse policies. *arXiv preprint arXiv:1906.00088*, 2019.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Smith *et al.*, 2018] Matthew Smith, Herke Hoof, and Joelle Pineau. An inference-based policy gradient method for learning options. In *International Conference on Machine Learning*, pages 4710–4719, 2018.
- [Sohrabi *et al.*, 2016] Shirin Sohrabi, Anton V Riabov, Octavian Udrea, and Oktie Hassanzadeh. Finding diverse high-quality plans for hypothesis generation. In *ECAI*, pages 1581–1582, 2016.
- [Sriperumbudur *et al.*, 2010] Bharath K Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Bernhard Schölkopf, and Gert RG Lanckriet. Non-parametric estimation of integral probability metrics. In *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, pages 1428–1432. IEEE, 2010.
- [Srivastava *et al.*, 2007] Biplav Srivastava, Tuan Anh Nguyen, Alfonso Gerevini, Subbarao Kambhampati, Minh Binh Do, and Ivan Serina. Domain independent approaches for finding diverse plans. In *IJCAI*, pages 2016–2022, 2007.
- [Sutton *et al.*, 2000] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
- [Thomas, 2015] Philip S Thomas. *Safe reinforcement learning*. PhD thesis, University of Massachusetts Libraries, 2015.
- [Toussaint and Lopes, 2017] Marc Toussaint and Manuel Lopes. Multi-bound tree search for logic-geometric programming in cooperative manipulation domains. In *(ICRA 2017)*, 2017.