
Represent Your Own Policies: Learning with Policy-extended Value Function Approximator

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 We study Policy-extended Value Function Approximator (PeVFA) in Reinforce-
2 ment Learning (RL), which extends conventional value function approximator
3 (VFA) to take as input not only the state (and action) but also an explicit policy
4 representation. Such an extension enables PeVFA to preserve values of multi-
5 ple policies at the same time and brings an appealing characteristic, i.e., *value*
6 *generalization among policies*. We formally analyze the value generalization under
7 Generalized Policy Iteration (GPI). From theoretical and empirical lens, we
8 show that generalized value estimates offered by PeVFA may have lower initial
9 approximation error to true values of successive policies, which is expected to
10 improve consecutive value approximation during GPI. Based on above clues, we
11 introduce a new form of GPI with PeVFA which leverages the value generalization
12 along policy improvement path. Moreover, we propose a representation learning
13 framework for RL policy, providing several approaches to learn effective policy em-
14 beddings from policy network parameters or state-action pairs. In our experiments,
15 we evaluate the efficacy of value generalization offered by PeVFA and policy
16 representation learning in several OpenAI Gym continuous control tasks. For a
17 representative instance of algorithm implementation, Proximal Policy Optimization
18 (PPO) re-implemented under the paradigm of GPI with PeVFA achieves about 40%
19 performance improvement on its vanilla counterpart in most environments.

20 1 Introduction

21 Reinforcement Learning (RL) has been widely considered as a promising way to learn optimal
22 policies in many decision-making problems [35, 31, 53, 65, 47, 62, 16]. One fundamental element of
23 RL is value function which defines the long-term evaluation of a policy. With function approximation
24 (e.g., deep neural networks), a value function approximator (VFA) is able to approximate the values
25 of a policy under large and continuous state spaces. As commonly recognized, most RL algorithms
26 can be described as Generalized Policy Iteration (GPI) [55]. As illustrated on the left of Figure 1,
27 at each iteration the VFA is trained to approximate the true values of current policy (i.e., policy
28 evaluation), regarding which the policy is further improved (i.e., policy improvement). The value
29 function approximation error hinders the effectiveness of policy improvement and then the overall
30 optimality of GPI [5, 46]. Unfortunately, such errors are inevitable under function approximation. A
31 large number of samples are usually required to ensure high-quality value estimates, resulting in the
32 sample-inefficiency of deep RL algorithms. Therefore, this raises an urgent need for more efficient
33 value approximation methods [61, 4, 12, 25].

34 An intuitive idea to improve the efficiency value approximation is to leverage the knowledge on
35 the values of previous encountered policies. However, a conventional VFA usually approximates
36 the values of one policy and values learned from old policies are over-written gradually during




Figure 1: Generalized Policy Iteration (GPI) with function approximation. *Left*: GPI with conventional value function approximator V_ϕ . *Right*: GPI with PeVFA $\mathbb{V}_\theta(\chi_\pi)$ (Sec. 3) where extra generalization steps exist. The subscripts of policy π and value function parameters ϕ, θ denote the iteration number. The squiggle lines represent non-perfect approximation of true values.

37 the learning process. This means that the previously learned knowledge cannot be preserved and
 38 utilized with one conventional VFA. Thus, such limitations prevent the potentials to leverage the
 39 previous knowledge for future learning. In this paper, we study Policy-extended Value Function
 40 Approximator (PeVFA), which additionally takes an explicit policy representation as input in contrast
 41 to conventional VFA. Thanks to the policy representation input, PeVFA is able to approximate values
 42 for multiple policies and induces value generalization among policies. We formally analyze the
 43 generalization of approximate values among policies in a general form. From both theoretical and
 44 empirical lens, we show that the generalized value estimates can be closer to the true values of
 45 the successive policy, which can be beneficial to consecutive value approximation along the policy
 46 improvement path, called *local generalization*. Based on above clues, we introduce a new form
 47 of GPI with PeVFA (the right of Figure 1) that leverages the local generalization to improve the
 48 efficiency of consecutive value approximation along the policy improvement path.

49 One key point of GPI with PeVFA is the representation of policy since it determines how PeVFA gen-
 50 eralizes the values. For this, we propose a framework to learn effective low-dimensional embedding
 51 of RL policy. We use network parameters or state-action pairs as policy data and encode them into
 52 low-dimensional embeddings; then the embeddings are trained to capture the effective information
 53 through contrastive learning and policy recovery. Finally, we evaluate the efficacy of GPI with PeVFA
 54 and our policy representations. In principle, GPI with PeVFA is general and can be implemented
 55 in different ways. As a practical instance, we re-implement Proximal Policy Optimization (PPO)
 56 with PeVFA and propose PPO-PeVFA algorithm. Our experimental results on several OpenAI Gym
 57 continuous control tasks demonstrate the effectiveness of both value generalization offered by PeVFA
 58 and learned policy representations, with an about 40% improvement in average returns achieved by
 59 our best variants on standard PPO in most tasks.

60 We summarize our main contributions below. 1) We study the value generalization among policies
 61 induced by PeVFA. From both theoretical and empirical aspects, we shed the light on the situations
 62 where the generalization can be beneficial to the learning along policy improvement path. 2) We
 63 propose a framework for policy representation learning. To our knowledge, we make the first attempt
 64 to learn a low-dimensional embedding of over 10k network parameters for an RL policy. 3) We
 65 introduce GPI with PeVFA that leverages the value generalization in a general form. Our experimental
 66 results demonstrate the potential of PeVFA in deriving practical and more effective RL algorithms.

67 2 Related Work

68 **Extensions of Conventional Value Function.** Sutton et al. [56] propose General Value Functions
 69 (GVFs) as a general form of knowledge representation of rewards and arbitrary cumulants. Later,
 70 conventional value functions are extended to take extra inputs for different purposes of generalization.
 71 One notable work is Universal Value Function Approximator (UVFA) [45], which is proposed to
 72 generalize values among different goals for goal-conditioned RL. UVFA is further developed in
 73 [1, 37, 9] and influences the occurrence of other value function extensions in context-based Meta-RL
 74 [43, 29], Hierarchical RL [64] and multiagent RL [19, 14] and etc. Most of the above works study
 75 how to generalize the policy or value function among extrinsic factors, i.e., environments, tasks and
 76 opponents; while we mainly study the value generalization among policies along policy improvement
 77 path, an intrinsic learning process of the agent itself.

78 **Policy Embedding and Representation.** Although not well studied, representation (or embedding)
79 learning for RL policies is involved in a few works [18, 14, 3]. The most common way to learn a
80 policy representation is to extract from interaction experiences. As a representative, Grover et al. [14]
81 propose learning the representation of opponent policy from interaction trajectories with a generative
82 policy recovery loss and a discriminative triplet loss. These losses are later adopted in [64, 42].
83 Another straightforward idea is to represent policy parameters. Network Fingerprint [17] is such a
84 differentiable representation that uses the concatenation of the vectors of action distribution outputted
85 by policy network on a set of probing states. The probing state set is co-optimized along with the
86 primary learning objective, which can be non-trivial especially when the dimensionality of the set is
87 high. Besides, some early attempts in learning low-dimensional embedding of policy parameters are
88 studies in Evolutionary Algorithms [13, 44], mainly with the help of VAE [23]. Our work introduce a
89 learning framework of policy representation including both above two perspectives.

90 **PVN and PVFs.** Recently, several works study the generalization among policy space. Harb et al.
91 [17] propose Policy Evaluation Network (PVN) to directly approximate the distribution of policy
92 π 's objective function $J(\pi) = \mathbb{E}_{\rho_0}[v^\pi(s_0)]$ with initial state $s_0 \sim \rho_0$. PVN takes as input Network
93 Fingerprint (mentioned above) of policy network. After training on a pre-collected set of policies, a
94 random initialized policy can be optimized in a zero-shot manner with the policy gradients of PVN by
95 backpropagating through the differentiable policy input. We call such gradients *GTPI* for short below.
96 Similar ideas are later integrated with task-specific context learning in multi-task RL [42], leveraging
97 the generalization among policies and tasks for fast policy adaptation on new tasks. In PVN [17],
98 as an early attempt, the generalization among policies is studied with small policy network and
99 simple tasks; besides, the most regular online learning setting is not studied. Concurrent to our work,
100 Faccio and Schmidhuber [10] propose a class of Parameter-based Value Functions (PVFs) that take
101 vectorized policy parameters as inputs. Based on PVFs, new policy gradient algorithms are introduced
102 in the form of a combination of conventional policy gradients and GTPI (i.e., by backpropagating
103 through policy parameters in PVFs). Except for zero-shot policy optimization as conducted in PVN,
104 PVFs are also evaluated for online policy learning. Due to directly taking parameters as input, PVFs
105 suffer from the curse of dimensionality when the number of parameters is high. Besides, GTPI can
106 be non-trivial to rein since policy parameter space are complex and extrapolation generalization
107 error can be large when the value function is only trained on finite policies (usually much fewer than
108 state-action samples) thus further resulting in erroneous policy gradients.

109 Our work differs with PVFs from several aspects. First, we make use of learned policy representation
110 rather than policy network parameters. Second, we do not resort to GTPI for the policy update
111 in our algorithms but focus on utilizing value generalization for more efficient value estimation in
112 GPI. Furthermore, we shed the light on two important problems — how value generalization among
113 policies can happen formally and whether it is beneficial to learning or not — which are neglected in
114 previous works from both theoretical and empirical lens.

115 3 Policy-extended Value Function Approximator

116 In this section, we propose Policy-extended Value Function Approximator (PeVFA), an extension
117 of conventional VFA that explicitly takes as input a policy representation. First, we introduce the
118 formulation (Sec. 3.1), then we study value generalization among policies theoretically (Sec. 3.2)
119 along with some empirical evidences (Sec. 3.3). Finally, we derive a new form of GPI (Sec. 3.4).

120 3.1 Formulation

121 Consider a Markov Decision Process (MDP) defined as $\langle \mathcal{S}, \mathcal{A}, r, \mathcal{P}, \gamma \rangle$ where \mathcal{S} is the state space, \mathcal{A}
122 is the action space, r is the (bounded) reward function, \mathcal{P} is the transition function and $\gamma \in [0, 1)$ is
123 the discount factor. A policy $\pi \in P(\mathcal{A})^{|\mathcal{S}|}$ defines the distribution over all actions for each state. The
124 goal of an RL agent is to find an optimal policy π^* that maximizes the expected long-term discounted
125 return. The state-value function $v^\pi(s)$ is defined as the expected discounted return obtained through
126 following the policy π from a state s : $v^\pi(s) = \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r_{t+1} | s_0 = s]$ for where $r_{t+1} = r(s_t, a_t)$.
127 We use V^π to denote the vectorized form of value function.

128 In a general form, we define *policy-extended value function* $\mathbb{V} : \mathcal{S} \times \Pi \rightarrow \mathbb{R}$ over state and policy
129 space: $\mathbb{V}(s, \pi) = v^\pi(s)$ for all $s \in \mathcal{S}$ and $\pi \in \Pi$. In this paper, we focus on $\mathbb{V}(s, \pi)$ and policy-
130 extended action-value function $\mathbb{Q}(s, a, \pi)$ can be obtained similarly. We use $\mathbb{V}(\pi)$ to denote the value

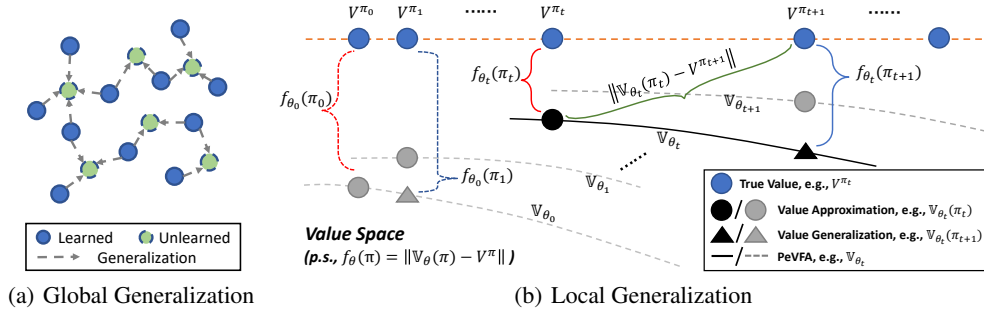


Figure 2: Illustrations of value generalization among policies of PeVFA. Each circle denotes value function (estimate) of a policy. (a) *Global Generalization*: values learned from known policies can be generalized to unknown policies. (b) *Local Generalization*: values of previous policies (e.g., π_t) can be generalized to successive policies (e.g., π_{t+1}) along policy improvement path.

131 vector for all states in the following. The key point is that PeVFA \mathbb{V} is able to preserve the values of
 132 multiple policies. With function approximation, a PeVFA is expected to approximate the values of
 133 policies among policy space, i.e., $\{V^\pi\}_{\pi \in \Pi}$ and then enable value generalization among policies.

134 Formally, given a function $g : \Pi \rightarrow \mathcal{X} \subseteq \mathbb{R}^n$ that maps any policy π to an n -dimensional representation
 135 $\chi_\pi = g(\pi) \in \mathcal{X}$, a PeVFA \mathbb{V}_θ with parameter $\theta \in \Theta$ is to minimize the approximation error
 136 over all possible states and policies generally:

$$F_{\mu, \rho, \rho}(\theta, g, \Pi) = \sum_{\pi \in \Pi} \mu(\pi) \|\mathbb{V}_\theta(\chi_\pi) - V^\pi\|_{p, \rho}, \quad (1)$$

137 where μ, ρ are distributions over policies and states respectively, $\|f\|_{p, \rho} = (\int_{\mathcal{S}} \rho(ds) |f(s)|^p)^{1/p}$ is
 138 ρ -weighted L_p -norm [26, 46] for any $f : \mathcal{S} \rightarrow \mathbb{R}$. The policy distribution μ of interest depends on
 139 the scenario where value generalization is considered. As illustrated in Figure 2, we provide two
 140 value generalization scenarios. In the global generalization scenario, a uniform distribution over
 141 known policy set may be considered with a general purpose of value generalization for unknown
 142 policies. For the specific local generalization scenario along policy improvement path during GPI,
 143 a sophisticated distribution that adaptively weights recent policies more during the learning process
 144 may be more suitable in this case. In the following, we care more about the local generalization
 145 scenario and use uniform state distribution ρ and L_2 -norm for demonstration. The subscripts are
 146 omitted and we use $\|\cdot\|$ for clarity.

147 3.2 Theoretical Analysis on Value Generalization among Policies

148 In this part, we theoretically analyze the value generalization among policies induced by PeVFA. We
 149 start from a two-policy case and study whether the value approximation learned for one policy can be
 150 generalized to the other one. Later, we study the local generalization scenario (Figure 2(b)) and shed
 151 the light on the superiority of PeVFA for GPI. All the proofs are provided in Appendix A.

152 For the convenience of demonstration, we use an identical policy representation function, i.e., $\chi_\pi = \pi$,
 153 and define the approximation loss of PeVFA \mathbb{V}_θ for any policy $\pi \in \Pi$ as $f_\theta(\pi) = \|\mathbb{V}_\theta(\pi) - V^\pi\| \geq 0$.
 154 We use the following definitions for a formal description of value approximation process with PeVFA
 155 and local property of loss function f_θ that influences generalization [40, 63] respectively:

156 **Definition 1 (π -Value Approximation)** We define a value approximation process $\mathcal{P}_\pi : \Theta \rightarrow \Theta$
 157 with PeVFA as a γ -contraction mapping on the approximation loss for policy π , i.e., for $\hat{\theta} = \mathcal{P}_\pi(\theta)$,
 158 we have $f_{\hat{\theta}}(\pi) \leq \gamma f_\theta(\pi)$ where $\gamma \in [0, 1)$.

159 **Definition 2 (L -Continuity)** We call f_θ is L -continuous at policy π if f_θ is Lipschitz continuous at
 160 π with a constant $L \in [0, \infty)$, i.e., $|f_\theta(\pi) - f_\theta(\pi')| \leq L \cdot d(\pi, \pi')$ for $\pi' \in \Pi$ with some distance
 161 metric d for policy space Π .

162 With Definition 1, the consecutive value approximation for the policies along policy improvement path
 163 during GPI can be described as: $\theta_{-1} \xrightarrow{\mathcal{P}_{\pi_0}} \theta_0 \xrightarrow{\mathcal{P}_{\pi_1}} \theta_1 \xrightarrow{\mathcal{P}_{\pi_2}} \dots$, as the green arrows illustrated in
 164 Figure 1. One may refer to Appendix A.1 for a discussion on the rationality of the two definitions.

165 To start our analysis, we first study the generalized value approximation loss in a two-policy case
 166 where only the value of policy π_1 is approximated by PeVFA as below:

167 **Lemma 1** For $\theta \xrightarrow{\mathcal{P}_{\pi_1}} \hat{\theta}$, if $f_{\hat{\theta}}$ is \hat{L} -continuous at π_1 and $f_{\theta}(\pi_1) \leq f_{\theta}(\pi_2)$, we have: $f_{\hat{\theta}}(\pi_2) \leq$
 168 $\gamma f_{\theta}(\pi_2) + \mathcal{M}(\pi_1, \pi_2, \hat{L})$, where $\mathcal{M}(\pi_1, \pi_2, \hat{L}) = \hat{L} \cdot d(\pi_1, \pi_2)$.

169 **Corollary 1** \mathcal{P}_{π_1} is γ_g -contraction ($\gamma_g \in [0, 1)$) for π_2 when $f_{\theta}(\pi_2) > \frac{\hat{L} \cdot d(\pi_1, \pi_2)}{1 - \gamma}$.

170 Lemma 1 shows that the post- \mathcal{P}_{π_1} approximation loss for π_2 is upper bounded by a generalized
 171 contraction of prior loss plus a locality margin term \mathcal{M} which is related to π_1 , π_2 and the locality
 172 property of $f_{\hat{\theta}}$. In general, the form of \mathcal{M} depends on the local property assumed. Some higher-
 173 order variants are provided in Appendix A.2. For a step further, Corollary 1 reveals the condition
 174 where a contraction on value approximation loss for π_2 is achieved when PeVFA is only trained to
 175 approximate the values of π_1 . Concretely, such a condition is apt to reach with tighter contraction for
 176 policy π_1 is, closer two policies, or smoother approximation loss function $f_{\hat{\theta}}$.

177 Then we consider the local generalization scenario as illustrated in Figure 2(b). For any iteration t
 178 of GPI, the values of current policy π_t are approximated by PeVFA, followed by a improved policy
 179 π_{t+1} whose values are to be approximated in the next iteration. The value generalization from each
 180 π_t and π_{t+1} can be similarly considered as the two-policy case. In addition to the former results, we
 181 shed the light on the value generalization loss of PeVFA along policy improvement path below:

182 **Lemma 2** For $\theta_{-1} \xrightarrow{\mathcal{P}_{\pi_0}} \theta_0 \xrightarrow{\mathcal{P}_{\pi_1}} \theta_1 \xrightarrow{\mathcal{P}_{\pi_2}} \dots$ with γ_t for each \mathcal{P}_{π_t} , if f_{θ_t} is L_t -continuous at π_t
 183 for any $t \geq 0$, we have $f_{\theta_t}(\pi_{t+1}) \leq \gamma_t f_{\theta_{t-1}}(\pi_t) + \mathcal{M}_t$, where $\mathcal{M}_t = L_t \cdot d(\pi_t, \pi_{t+1})$.

184 **Corollary 2** By induction, we have $f_{\theta_t}(\pi_{t+1}) \leq \prod_{i=0}^t \gamma_i f_{\theta_{-1}}(\pi_0) + \sum_{i=0}^{t-1} \prod_{j=i+1}^t \gamma_j \mathcal{M}_i + \mathcal{M}_t$.

185 The above results indicate that the value generalization loss can be recursively bounded and has
 186 a upper bound formed by a repeated contraction on initial loss plus the accumulation of locality
 187 margins induced from each local generalization. An infinity-case discussion for Corollary 2 is in
 188 Appendix A.5. The next question is whether PeVFA with value generalization among policies is
 189 preferable to the conventional VFA. To this end, we introduce a desirable condition which reveals the
 190 superiority of PeVFA during consecutive value approximation along the policy improvement path:

191 **Theorem 1** During $\theta_{-1} \xrightarrow{\mathcal{P}_{\pi_0}} \theta_0 \xrightarrow{\mathcal{P}_{\pi_1}} \theta_1 \xrightarrow{\mathcal{P}_{\pi_2}} \dots$, for any $t \geq 0$, if $f_{\theta_t}(\pi_t) + f_{\theta_t}(\pi_{t+1}) \leq$
 192 $\|V^{\pi_t} - V^{\pi_{t+1}}\|$, then $f_{\theta_t}(\pi_{t+1}) \leq \|\mathbb{V}_{\theta_t}(\pi_t) - V^{\pi_{t+1}}\|$.

193 Theorem 1 shows that the generalized value estimates $\mathbb{V}_{\theta_t}(\pi_{t+1})$ can be closer to the true values of
 194 policy π_{t+1} than $\mathbb{V}_{\theta_t}(\pi_t)$. Note that $\mathbb{V}_{\theta_t}(\pi_t)$ is the value approximation for π_t which is equivalent
 195 to the counterpart V_{ϕ_t} for a conventional VFA as value generalization among policies does not
 196 exist. To consecutive value approximation along policy improvement path, this means that the value
 197 generalization of PeVFA has the potential to offer closer start points at each iteration. If such closer
 198 start points can often exist, we expect PeVFA to be preferable to conventional VFA since value
 199 approximation can be more efficient with PeVFA and it in turn facilitates the overall GPI process.

200 However, the condition in Theorem 1 is not necessarily met in practice. Intuitively, it depends on the
 201 locality margins that may be related to function family and optimization method of PeVFA, as well
 202 as the scale of policy improvement. We leave these further theoretical investigations for future work.
 203 Instead, we empirically examine the existence of such desirable generalizations in the following.

204 3.3 Empirical Evidences

205 We empirically investigate the value generalization of PeVFA with didactic environments. In this
 206 section, PeVFA \mathbb{V}_{θ} is parameterized by neural network and we use the concatenation of all weights
 207 and biases of the policy network as a straightforward representation χ_{π} for each policy, called *Raw*
 208 *Policy Representation (RPR)*. Experimental details are provided in Appendix B.

209 First, we demonstrate the global generalization (illustrated in Figure 2(a)) in a continuous 2D Point
 210 Walker environment. We build the policy set Π with synthetic policies, each of which is a randomly
 211 initialized 2-layer *tanh*-activated neural network with 2 units for each layer. The size of Π is 20k and
 212 the behavioral diversity of synthetic policies is verified (see Figure 7(b) in Appendix). We divide Π
 213 into training set (i.e., known policies Π_0) and testing set (i.e., unseen policies Π_1). We rollout the




Figure 3: Empirical evidences of two kinds of generalization of PeVFA. (a) *Global generalization*: PeVFA shows comparable value estimation performance on testing policy set (red) after learning on training policy set (blue). (b) *Local generalization*: PeVFA ($\mathbb{V}_\theta(\chi_\pi)$) shows lower losses than conventional VFA (V_ϕ) before and after the value approximation training for successive policies along policy improvement path. In (b), the left axis is for approximation loss (lower is better) and the right axis is for average return as a reference of the policy learning process (green curve).

214 policies in the environment to collect trajectories, based on which we perform value approximation
 215 training. Our results show that a PeVFA trained on Π_0 achieves reasonable generalization performance
 216 when evaluating on Π_1 . The average losses on training and testing set are 1.782 and 2.071 over 6
 217 trials. Figure 3(a) shows the value predictions for policies from training and testing set (100 for each).

218 Next, we investigate the value generalization along policy improvement path, i.e., local generalization
 219 as in Figure 2(b). We use a 2-layer 8-unit policy network trained by standard PPO algorithm [50] in
 220 MuJoCo continuous control tasks. Parallel to the conventional value network $V_\phi(s)$ (i.e., VFA) in
 221 PPO, we set a PeVFA network $\mathbb{V}_\theta(s, \chi_\pi)$ as a reference for the comparison on value approximation
 222 loss. Compared to V_ϕ , PeVFA $\mathbb{V}_\theta(s, \chi_\pi)$ takes RPR as input and approximates the values of all
 223 historical policies ($\{\pi_i\}_{i=0}^t$) in addition. We compare the value approximation losses of V_ϕ (red) and
 224 \mathbb{V}_θ (blue) before (solid) and after (dashed) updating with on-policy samples collected by the improved
 225 policy π_{t+1} at each iteration. Figure 3(b) shows the results for InvertedPendulum-v1 and Ant-v1.
 226 Results for all 7 MuJoCo tasks can be found in Appendix B.2. By comparing approximation losses
 227 before updating (red and blue solid curves), we can observe that the approximation loss of $\mathbb{V}_\theta(\chi_{\pi_{t+1}})$
 228 is almost consistently lower than that of V_{ϕ_t} . This means that the generalized value estimates
 229 offered by PeVFA are usually closer to the true values of π_{t+1} , demonstrating the consequence
 230 arrived in Theorem 1. For the dashed curves, it shows that PeVFA $\mathbb{V}_{\theta_{t+1}}(\chi_{\pi_{t+1}})$ can achieve lower
 231 approximation loss for π_{t+1} than conventional VFA $V_{\phi_{t+1}}$ after the same number of training with the
 232 same on-policy samples. The empirical evidence above indicates that PeVFA can be preferable to
 233 the conventional VFA for consecutive value approximation. The generalized value estimates along
 234 policy improvement path have the potential to expedite the process of GPI.

235 3.4 Reinforcement Learning with PeVFA

236 Based on the results above, we expect to leverage the value generalization of PeVFA to facilitate
 237 RL. In Algorithm 1, we propose a general description of RL algorithm under the paradigm of
 238 GPI with PeVFA. For each iteration, the interaction experiences of current policy and the policy

Algorithm 1 RL under the paradigm of GPI with PeVFA ($\mathbb{V}(s, \chi_\pi)$ is used for demonstration)

- 1: Initialize policy π_0 , policy representation model g , PeVFA \mathbb{V}_{-1} and experience buffer \mathcal{D}
 - 2: **for** iteration $t = 0, 1, \dots$ **do**
 - 3: Rollout policy π_t in the environment and obtain k trajectories $\mathcal{T}_t = \{\tau_i\}_{i=0}^k$
 - 4: Get representation $\chi_{\pi_t} = g(\pi)$ for policy π_t and add experiences $(\chi_{\pi_t}, \mathcal{T}_t)$ in buffer \mathcal{D}
 - 5: **if** $t \% M = 0$ **then**
 - 6: Update PeVFA $\mathbb{V}_{t-1}(s, \chi_{\pi_t})$ for previous policies with data $\{(\chi_{\pi_i}, \mathcal{T}_i)\}_{i=0}^{t-1}$
 - 7: Update policy representation model g , e.g., with approaches provided in Sec. 4
 - 8: **end if**
 - 9: Update PeVFA $\mathbb{V}_{t-1}(s, \chi_{\pi_t})$ for current policy χ_{π_t} and set $\mathbb{V}_t \leftarrow \mathbb{V}_{t-1}$
 - 10: Update π_t w.r.t $\mathbb{V}_t(s, \chi_{\pi_t})$ by policy improvement algorithm and set $\pi_{t+1} \leftarrow \pi_t$
 - 11: **end for**
-

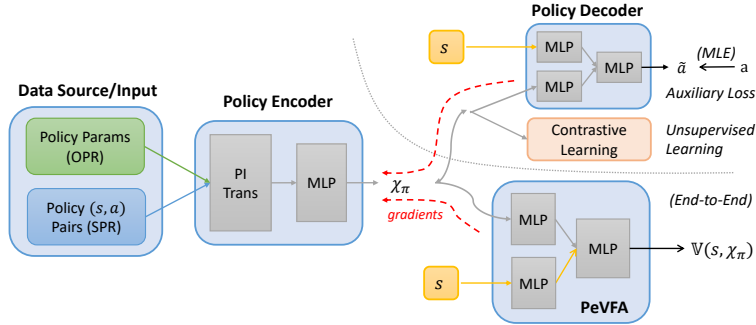


Figure 4: The framework of policy representation training. Policy network parameters used for OPR or policy state-action pairs used for SPR are fed into policy encoder with permutation-invariant (PI) transformations followed by an MLP, producing the representation χ_π . Afterwards, χ_π can be trained by gradients from the value approximation loss of PeVFA (i.e., End-to-End), as well as (optionally) the auxiliary loss of policy recovery or the contrastive learning (i.e., InfoNCE) loss.

239 representation are stored in a buffer (line 3-4). At an interval of M iterations, PeVFA is trained via
 240 value approximation for previous policies with the stored data and the policy representation model
 241 is updated according to the method used (line 5-8). This part is unique to PeVFA for preservation
 242 and generalization of knowledge of historical policies. Next, value approximation for current policy
 243 is performed with PeVFA (line 9). A key difference here is that the generalized value estimates
 244 (i.e., $\mathbb{V}_{t-1}(\chi_{\pi_t})$) are used as start points. Afterwards, a successive policy is obtained from typical
 245 policy improvement (line 10). Algorithm 1 can be implemented in different ways and we propose an
 246 instance implemented based on PPO [50] in our experiments later. In the next section, we introduce
 247 our methods for policy representation learning.

248 4 Policy Representation Learning

249 To derive practical deep RL algorithms, one key point is policy representation, i.e., a low-dimensional
 250 embedding of RL policy. Intuitively, policy representation influences the approximation and gener-
 251 alization of PeVFA. Thus, it is of interest to find an effective policy representation based on which
 252 the superiority of PeVFA can be leveraged to improve RL algorithms. To our knowledge, policy
 253 representation is not well studied and it remains unclear on how to obtain an effective representation
 254 for an RL policy in a general case in practice. In previous section, we demonstrate the effectiveness
 255 of using policy parameters as a naive representation when policy network is small, called RPR.
 256 However, a usual policy network may have large number of parameters, thus making it inefficient
 257 and even irrational to use RPR for approximation and generalization [17, 10]. More generally, policy
 258 parameters of the policy we wish to represent may not be accessible.

259 To this end, we propose a general framework of policy representation learning as illustrated in Figure
 260 4. The first thing to consider is data source, i.e., from which we can extract the information for an
 261 effective policy representation. Recall that the policy is a distribution over state and action space
 262 of high dimensionality. The features of such a distribution is not directly available. Therefore, we
 263 consider two kinds of data source below that indirectly contains the information of policies: 1) *Surface*
 264 *Policy Representation (SPR)*: The first data source is state-action pairs (or trajectories [14]), since
 265 they reflect how policy may behave under such states. This data source is general since no explicit
 266 form of policy is assumed. In a geometric view, learning policy representation from state-action pairs
 267 can be viewed as capturing the features of policy via scattering sample points on the curved surface
 268 of policy distribution. 2) *Origin Policy Representation (OPR)*: The other data source is parameters
 269 of policy since they determine the underlying form of policy distribution. Such a data source is often
 270 available during the learning process of deep RL algorithms when policy is parameterized by neural
 271 networks. Generally, we consider a policy network to be an MLP with well represented state features
 272 (e.g., features extracted by CNN for pixels or by LSTM for sequences) as input.

273 The remaining question is how we extract the policy representation from the data sources mentioned
 274 above. As shown in Figure 4, we use permutation-invariant (PI) transformations followed by an
 275 MLP to encode the data of policy π into an embedding χ_π for both SPR and OPR. For SPR, each

276 state-action pair of $\{(s_i, a_i)\}_{i=1}^k$ is fed into a common MLP, followed by a Mean-Reduce operation
 277 on the outputted features across k . For OPR, we perform PI transformation (similar as done for
 278 state-action pairs) inner-layer weights and biases $\{(w_i, b_i)\}_{i=1}^h$ for each layer first, where h denotes
 279 the number of nodes in this layer and w_i, b_i is the income weight vector from previous layer and
 280 the bias of i th node; then we concatenate encoding of layers and obtain the OPR. A illustrative
 281 description for the encoding of OPR is in Figure 12 of Appendix.

282 To train the policy embedding χ_π obtained above, the most straightforward way is to backpropagate
 283 the value approximation loss of PeVFA in an *End-to-End (E2E)* fashion as illustrated on the lower-
 284 right of Figure 4. In addition, we provide two self-supervised training losses for both OPR and SPR,
 285 as illustrated on the upper-right of Figure 4. The first one is an *auxiliary loss (AUX) of policy recovery*
 286 [14], i.e., to recover the action distributions of π from χ_π under different states. To be specific,
 287 an auxiliary policy decoder $\bar{\pi}(\cdot|s, \chi_\pi)$ is trained through behavioral cloning, formally to minimize
 288 cross-entropy objective $\mathcal{L}_{\text{AUX}} = -\mathbb{E}_{(s,a)} [\log \bar{\pi}(a|s, \chi_\pi)]$. For the second one, we propose to train
 289 χ_π by *Contrastive Learning (CL)* [54, 51]: policies are encouraged to be close to similar ones (i.e.,
 290 positive samples π^+), and to be apart from different ones (i.e., negative samples π^-) in representation
 291 space. For each policy, we construct positive samples by data augmentation on policy data, depending
 292 on SPR or OPR considered; and different policies along the policy improvement path naturally
 293 provide negative samples for each other. Finally, the embedding χ_π is optimized through minimizing
 294 the InfoNCE loss [41] below: $\mathcal{L}_{\text{CL}} = -\mathbb{E}_{(\pi^+, \{\pi^-\})} \left[\log \frac{\exp(\chi_\pi^T W \chi_{\pi^+})}{\exp(\chi_\pi^T W \chi_{\pi^+}) + \sum_{\pi^-} \exp(\chi_\pi^T W \chi_{\pi^-})} \right]$.

295 Now, the training of policy representation model in Algorithm 1 can be performed with any com-
 296 bination of data sources and training losses provided above. A pseudo-code of the overall policy
 297 representation training framework and complete implementation details are provided in Appendix D.

298 5 Experiments

299 In this section, we conduct experimental study with focus on the following questions:

300 **Question 1** *Can value generalization offered by PeVFA improve a deep RL algorithm in practice?*

301 **Question 2** *Can our proposed framework to learn effective policy representation?*

302 Our experiments are conducted in several OpenAI Gym continuous control tasks (one from Box2D
 303 and five from MuJoCo) [6, 58]. All experimental details and curves can be found in Appendix B.

304 **Algorithm Implementation.** We use PPO [50] as the basic algorithm and propose a representative
 305 implementation of Algorithm 1, called **PPO-PeVFA**. PPO is a policy optimization algorithm that
 306 follows the paradigm of GPI (Figure 1, left). A value network $V_\phi(s)$ with parameters ϕ (i.e.,
 307 conventional VFA) is trained to approximate the value of current policy π ; while π is optimized with
 308 respect to a surrogate objective [48] using advantages calculated by V_ϕ and GAE [49]. Compared with
 309 original PPO, PPO-PeVFA makes use of a PeVFA network $\mathbb{V}_\theta(s, \chi_\pi)$ with parameters θ rather than
 310 the conventional VFA $V_\phi(s)$, and follows the training scheme as in Algorithm 1. Note PPO-PeVFA
 311 uses the same policy optimization method as original PPO and only differs at value approximation.

312 **Baselines and Variants.** Except for original PPO as a default baseline, we use another two baselines:
 313 1) PPO-PeVFA with randomly generated policy representation for each policy, denoted by **Ran PR**;
 314 2) PPO-PeVFA with Raw Policy Representation (**RPR**), i.e., use the vector of all parameters of policy
 315 network as representation as adopted in PVFs [10]. Our variants of PPO-PeVFA differ at the policy
 316 representation used. In total, we consider 6 variants denoted by the combination of the policy data
 317 choice (i.e., **OPR, SPR**) and representation principle choice (i.e., **E2E, CL, AUX**).

318 **Experimental Details.** For all baselines and variants, we use a normal-scale policy network with
 319 2 layers and 64 units for each layer, resulting in over 3k to 10k (e.g., Ant-v1) policy parameters
 320 depending on the environments. We do not assume the access to pre-collected policies. Thus the
 321 size of policy set increases from 1 (i.e., the initial policy) during the learning process, to about 1k to
 322 2 for a single trial. The dimensionality of all kinds of policy representation expect for RPR is set
 323 to 64. The buffer D maintains recent 200k steps of interaction experience and the policy data of
 324 corresponding policy. The number of interaction step of each trial is 1M for InvDouPend-v1 and
 325 LunarLander-v2, 4M for Ant-v1 and 2M for the others.

326 **Results.** The overall experimental results are summarized in Table 1. In Figure 5, we provide
 327 aggregated results across all environments expect for InvDouPend-v1 and LunarLander-v2 (since

Table 1: Average returns (\pm half a std) over 10 trials for algorithms. Each result is the maximum evaluation along the training process. Top two values for each environment are bold.

| Environments | Benchmarks | | | Origin Policy Representation (Ours) | | | Surface Policy Representation (Ours) | | |
|----------------|------------|--------|-------------------|-------------------------------------|-------------------------------------|-------------------------------------|--------------------------------------|-------------------------------------|-------------------------------------|
| | PPO | Ran PR | RPR | E2E | CL | AUX | E2E | CL | AUX |
| HalfCheetah-v1 | 2621 | 2470 | 2325 \pm 399.27 | 3171 \pm 427.63 | 3725 \pm 348.55 | 3175 \pm 517.52 | 2774 \pm 233.39 | 3349 \pm 341.42 | 3216 \pm 506.39 |
| Hopper-v1 | 1639 | 1226 | 1097 \pm 213.47 | 2085 \pm 310.91 | 2351 \pm 231.11 | 2214 \pm 360.78 | 2227 \pm 297.35 | 2392 \pm 263.93 | 2577 \pm 217.73 |
| Walker2d-v1 | 1505 | 1269 | 317 \pm 152.68 | 1856 \pm 305.51 | 2038 \pm 315.51 | 2044 \pm 316.32 | 1930.57 \pm 456.02 | 2203 \pm 381.95 | 1980 \pm 325.54 |
| Ant-v1 | 2835 | 2742 | 2143 \pm 406.64 | 3581 \pm 185.43 | 4019 \pm 162.47 | 3784 \pm 268.99 | 3173 \pm 184.75 | 3632 \pm 134.27 | 3397 \pm 200.03 |
| InvDouPend-v1 | 9344 | 9355 | 8856 \pm 551.90 | 9357 \pm 0.29 | 9355 \pm 0.64 | 9355 \pm 0.68 | 9355 \pm 0.89 | 9356 \pm 0.96 | 9355 \pm 1.42 |
| LunarLander-v2 | 219 | 226 | -22 \pm 35.08 | 238 \pm 3.37 | 239 \pm 3.70 | 234 \pm 3.47 | 236 \pm 3.13 | 234 \pm 3.13 | 235 \pm 5.70 |

most algorithms achieve near-optimal results), where all returns are normalized by the results of PPO in Table 1. Full learning curves are omitted and can be found in Appendix F.2.

To Question 1. From Table 1, we can find that both PPO-PeVFA w/ OPR (E2E) and PPO-PeVFA w/ SPR (E2E) outperforms PPO in all 6 tasks, and achieve over 20% improvement in Figure 5. This demonstrates the effectiveness of PeVFA. Moreover, the improvement is further enlarged (to about 40%) by CL and AUX for both OPR and SPR. This indicates that the superiority of PeVFA can be further utilized with better policy representation that offers a more suitable space for value generalization.

To Question 2. In Table 1, consistent degeneration is observed for PPO-PeVFA w/ Ran PR due to the negative effects on generalization caused by the randomness and disorder of policy representation. This phenomenon seems to be more severe for PPO-PeVFA w/ RPR due to the complexity of high-dimensional parameter space. In contrast, the improvement achieved by our proposed PPO-PeVFA variants shows that effective policy representation can be learned from policy parameters (OPR) and state-action pairs (SPR) though value approximation loss (i.e., E2E) and further improved when additional self-supervised representation learning is involved as CL and AUX. Overall, OPR slightly outperforms SPR as CL does over AUX. We hypothesize that it is due to the stochasticity of state-action pairs which serve as inputs of SPR and training samples for AUX. This reveals the space for future improvement. In addition, we visualize the learned representation in Figure 6. We can observe that policies from different trials are locally continuous and show different modes of embedding trajectories due to random initialization and optimization; while a global evolution among trials emerges with respect to policy performance.




Figure 5: Normalized averaged returns aggregated over 4 MuJoCo tasks.




Figure 6: A t-SNE visualization for representations learned by PPO-PeVFA OPR (E2E) in Ant-v1. In total, 6k policies from 5 trials (denoted by different markers) are plotted, which are colored according to average return.

6 Conclusion and Future Work

In this paper, we propose Policy-extended Value Function Approximator (PeVFA) and study value generalization among policies. We propose a new form of GPI based on PeVFA which is potentially preferable to conventional VFA for value approximation. Moreover, we propose a general framework to learn low-dimensional embedding of RL policy. Our experiments demonstrate the effectiveness of the generalization characteristic of PeVFA and our proposed policy representation learning methods.

Our work opens up some research directions on value generalization among policies and policy representation. A possible future study on the theory of value generalization among policies is to consider the interplay between approximation error, policy improvement and local generalization during GPI with PeVFA. Besides, analysis on influence factors of value generalization among policies (e.g., policy representation, architecture of PeVFA) and other utilization of PeVFA are expected. For better policy representation, inspirations on OPR may be got from studies on Manifold Hypothesis of neural network; the selection of more informative state-action pairs for SPR is also worth research.

374 **References**

- 375 [1] M. Andrychowicz, D. Crow, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel,
376 and W. Zaremba. Hindsight experience replay. In *NeurIPS*, pages 5048–5058, 2017.
- 377 [2] M. Andrychowicz, A. Raichuk, P. Stanczyk, M. Orsini, S. Girgin, R. Marinier, L. Hussenot, M. Geist,
378 O. Pietquin, M. Michalski, S. Gelly, and O. Bachem. What matters in on-policy reinforcement learning? A
379 large-scale empirical study. *CoRR*, abs/2006.05990, 2020.
- 380 [3] I. Arnekvist, D. Kragic, and J. A. Stork. VPE: variational policy embedding for transfer reinforcement
381 learning. In *ICRA*, pages 36–42, 2019.
- 382 [4] M. G. Bellemare, W. Dabney, and R. Munos. A distributional perspective on reinforcement learning. In
383 *ICML*, volume 70, pages 449–458, 2017.
- 384 [5] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-dynamic programming*, volume 3 of *Optimization and neural
385 computation series*. Athena Scientific, 1996.
- 386 [6] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym.
387 *CoRR*, abs/1606.01540, 2016.
- 388 [7] T. Chen, S. Kornblith, M. Norouzi, and G. E. Hinton. A simple framework for contrastive learning of
389 visual representations. *CoRR*, abs/2002.05709, 2020.
- 390 [8] L. Engstrom, A. Ilyas, S. Santurkar, D. Tsipras, F. Janoos, L. Rudolph, and A. Madry. Implementation
391 matters in deep RL: A case study on PPO and TRPO. In *ICLR*, 2020.
- 392 [9] B. Eysenbach, X. Geng, S. Levine, and R. R. Salakhutdinov. Rewriting history with inverse RL: hindsight
393 inference for policy improvement. In *NeurIPS*, 2020.
- 394 [10] F. Faccio and J. Schmidhuber. Parameter-based value functions. *CoRR*, abs/2006.09226, 2020.
- 395 [11] H. Fu, H. Tang, J. Hao, C. Chen, X. Feng, D. Li, and W. Liu. Towards effective context for meta-
396 reinforcement learning: an approach based on contrastive learning. *CoRR*, abs/2009.13891, 2020.
- 397 [12] S. Fujimoto, H. v. Hoof, and D. Meger. Addressing function approximation error in actor-critic methods.
398 In *ICML*, 2018.
- 399 [13] A. Gaier, A. Asteroth, and J. Mouret. Discovering representations for black-box optimization. In *GECCO*,
400 pages 103–111, 2020.
- 401 [14] A. Grover, M. Al-Shedivat, J. K. Gupta, Y. Burda, and H. Edwards. Learning policy representations in
402 multiagent systems. In *ICML*, volume 80, pages 1797–1806, 2018.
- 403 [15] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep
404 reinforcement learning with a stochastic actor. In *ICML*, pages 1856–1865, 2018.
- 405 [16] D. Hafner, T. P. Lillicrap, J. Ba, and M. Norouzi. Dream to control: Learning behaviors by latent
406 imagination. In *ICLR*, 2020.
- 407 [17] J. Harb, T. Schaul, D. Precup, and P. Bacon. Policy evaluation networks. *CoRR*, abs/2002.11833, 2020.
- 408 [18] K. Hausman, J. Tobias Springenberg, Z. Wang, N. Heess, and M. A. Riedmiller. Learning an embedding
409 space for transferable robot skills. In *ICLR*, 2018.
- 410 [19] H. He and J. L. Boyd-Graber. Opponent modeling in deep reinforcement learning. In *ICML*, volume 48,
411 pages 1804–1813, 2016.
- 412 [20] K. He, H. Fan, Y. Wu, S. Xie, and R. B. Girshick. Momentum contrast for unsupervised visual representation
413 learning. In *CVPR*, pages 9726–9735, 2020.
- 414 [21] M. Igl, G. Farquhar, J. Luketina, W. Boehmer, and S. Whiteson. The impact of non-stationarity on
415 generalisation in deep reinforcement learning. *CoRR*, abs/2006.05826, 2020.
- 416 [22] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang. On large-batch training for
417 deep learning: Generalization gap and sharp minima. In *ICLR*, 2017.
- 418 [23] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- 419 [24] I. Kostrikov, D. Yarats, and R. Fergus. Image augmentation is all you need: Regularizing deep reinforcement
420 learning from pixels. *CoRR*, abs/2004.13649, 2020.

- 421 [25] A. Kuznetsov, P. Shvechikov, A. Grishin, and D. P. Vetrov. Controlling overestimation bias with truncated
422 mixture of continuous distributional quantile critics. In *ICML*, volume 119, pages 5556–5566, 2020.
- 423 [26] M. G. Lagoudakis and R. Parr. Least-squares policy iteration. *J. Mach. Learn. Res.*, 4:1107–1149, 2003.
- 424 [27] Q. Lan, Y. Pan, A. Fyshe, and M. White. Maxmin q-learning: Controlling the estimation bias of q-learning.
425 In *ICLR*, 2020.
- 426 [28] M. Laskin, K. Lee, A. Stooke, L. Pinto, P. Abbeel, and A. Srinivas. Reinforcement learning with augmented
427 data. *CoRR*, abs/2004.14990, 2020.
- 428 [29] K. Lee, Y. Seo, S. Lee, H. Lee, and J. Shin. Context-aware dynamics model for generalization in
429 model-based reinforcement learning. *CoRR*, abs/2005.06800, 2020.
- 430 [30] A. Levy, G. D. Konidaris, R. P. Jr., and K. Saenko. Learning multi-level hierarchies with hindsight. In
431 *ICLR*, 2019.
- 432 [31] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous
433 control with deep reinforcement learning. In *ICLR*, 2015.
- 434 [32] H. Liu, K. Simonyan, and Y. Yang. DARTS: differentiable architecture search. In *ICLR*, 2019.
- 435 [33] R. Luo, F. Tian, T. Qin, E. Chen, and T. Liu. Neural architecture optimization. In *NeurIPS*, pages
436 7827–7838, 2018.
- 437 [34] L. V. D. Maaten and G. E. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:
438 2579–2605, 2008.
- 439 [35] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. A. Riedmiller,
440 A. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran,
441 D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*,
442 518(7540):529–533, 2015.
- 443 [36] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. P. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu.
444 Asynchronous methods for deep reinforcement learning. In *ICML*, 2016.
- 445 [37] O. Nachum, S. Gu, H. Lee, and S. Levine. Data-efficient hierarchical reinforcement learning. *CoRR*,
446 abs/1805.08296, 2018.
- 447 [38] O. Nachum, S. Gu, H. Lee, and S. Levine. Near-optimal representation learning for hierarchical reinforce-
448 ment learning. In *ICLR*, 2019.
- 449 [39] Y. E. Nesterov and B. T. Polyak. Cubic regularization of newton method and its global performance. *Math.*
450 *Program.*, 108(1):177–205, 2006.
- 451 [40] R. Novak, Y. Bahri, D. A. Abolafia, J. Pennington, and J. Sohl-Dickstein. Sensitivity and generalization in
452 neural networks: an empirical study. In *ICLR*, 2018.
- 453 [41] A. Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *CoRR*,
454 abs/1807.03748, 2018.
- 455 [42] R. Raileanu, M. Goldstein, A. Szlam, and R. Fergus. Fast adaptation via policy-dynamics value functions.
456 *CoRR*, abs/2007.02879, 2020.
- 457 [43] K. Rakelly, A. Zhou, C. Finn, S. Levine, and D. Quillen. Efficient off-policy meta-reinforcement learning
458 via probabilistic context variables. In *ICML*, volume 97, pages 5331–5340, 2019.
- 459 [44] N. Rakicevic, A. Cully, and P. Kormushev. Policy manifold search: Exploring the manifold hypothesis for
460 diversity-based neuroevolution. *CoRR*, abs/2104.13424, 2021.
- 461 [45] T. Schaul, D. Horgan, K. Gregor, and D. Silver. Universal value function approximators. In *ICML*,
462 volume 37, pages 1312–1320, 2015.
- 463 [46] B. Scherrer, M. Ghavamzadeh, V. Gabillon, B. Lesner, and M. Geist. Approximate modified policy iteration
464 and its application to the game of tetris. *J. Mach. Learn. Res.*, 16:1629–1676, 2015.
- 465 [47] J. S Schreck, C. W Coley, and K. JM Bishop. Learning retrosynthetic planning through simulated
466 experience. *ACS central science*, 5(6):970–981, 2019.

- 467 [48] J. Schulman, S. Levine, P. Abbeel, M. I. Jordan, and P. Moritz. Trust region policy optimization. In *ICML*,
468 pages 1889–1897, 2015.
- 469 [49] J. Schulman, P. Moritz, S. Levine, M. I. Jordan, and P. Abbeel. High-dimensional continuous control using
470 generalized advantage estimation. In *ICLR*, 2016.
- 471 [50] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms.
472 *CoRR*, abs/1707.06347, 2017.
- 473 [51] M. Schwarzer, A. Anand, R. Goel, R. D. Hjelm, A. C. Courville, and P. Bachman. Data-efficient
474 reinforcement learning with momentum predictive representations. *CoRR*, abs/2007.05929, 2020.
- 475 [52] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. A. Riedmiller. Deterministic policy gradient
476 algorithms. In *ICML*, pages 387–395, 2014.
- 477 [53] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser,
478 I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner,
479 I. Sutskever, T. P. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game
480 of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- 481 [54] A. Srinivas, M. Laskin, and P. Abbeel. CURL: contrastive unsupervised representations for reinforcement
482 learning. *CoRR*, abs/2004.04136, 2020.
- 483 [55] R. S. Sutton and A. G. Barto. *Reinforcement learning - an introduction*. Adaptive computation and machine
484 learning. MIT Press, 1998.
- 485 [56] R. S. Sutton, J. Modayil, M. Delp, T. Degris, P. M. Pilarski, A. White, and D. Precup. Horde: a scalable
486 real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *AAMAS*,
487 pages 761–768, 2011.
- 488 [57] A. Tacchetti, H. F. Song, P. A. M. Mediano, V. Flores Zambaldi, J. Kramár, N. C. Rabinowitz, Th. Graepel,
489 M. Botvinick, and P. W. Battaglia. Relational forward models for multi-agent learning. In *ICLR*, 2019.
- 490 [58] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. *2012 IEEE/RSJ*
491 *International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012.
- 492 [59] Y. H. Tsai, Y. Wu, R. Salakhutdinov, and L. Morency. Demystifying self-supervised learning: An
493 information-theoretical framework. *CoRR*, abs/2006.05576, 2020.
- 494 [60] T. Unterthiner, D. Keysers, S. Gelly, O. Bousquet, and I. O. Tolstikhin. Predicting neural network accuracy
495 from weights. *CoRR*, abs/2002.11448, 2020.
- 496 [61] H. v. Hasselt. Double q-learning. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and
497 A. Culotta, editors, *NeurIPS*, pages 2613–2621, 2010.
- 498 [62] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell,
499 T. Ewalds, P. Georgiev, J. Oh, D. Horgan, M. Kroiss, I. Danihelka, A. Huang, L. Sifre, T. Cai, J. P. Agapiou,
500 M. Jaderberg, A. S. Vezhnevets, R. Leblond, T. Pohlen, V. Dalibard, D. Budden, Y. Sulsky, J. Molloy,
501 T. L. Paine, C. Gulcehre, Z. Wang, T. Pfaff, Y. Wu, R. Ring, D. Yogatama, D. Wunsch, K. McKinney,
502 O. Smith, T. Schaul, T. Lillicrap, K. Kavukcuoglu, D. Hassabis, C. Apps, and D. Silver. Grandmaster level
503 in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- 504 [63] H. Wang, N. S. Keskar, C. Xiong, and R. Socher. Identifying generalization properties in neural networks.
505 *CoRR*, abs/1809.07402, 2018.
- 506 [64] R. Wang, R. Yu, B. An, and Z. Rabinovich. I²hrl: Interactive influence-based hierarchical reinforcement
507 learning. In *IJCAI*, pages 3131–3138, 2020.
- 508 [65] J. You, B. Liu, Z. Ying, V. S. Pande, and J. Leskovec. Graph convolutional policy network for goal-directed
509 molecular graph generation. In *NeurIPS 2018*, pages 6412–6422, 2018.

510 Checklist

- 511 1. For all authors...
- 512 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contribu-
513 tions and scope? [Yes]
- 514 (b) Did you describe the limitations of your work? [Yes] See the future work in Sec. 6.
- 515 (c) Did you discuss any potential negative societal impacts of your work? [No] Our work is on
516 general Reinforcement Learning study. No specific practical application is considered.
- 517 (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
- 518 2. If you are including theoretical results...
- 519 (a) Did you state the full set of assumptions of all theoretical results? [Yes]
- 520 (b) Did you include complete proofs of all theoretical results? [Yes]
- 521 3. If you ran experiments...
- 522 (a) Did you include the code, data, and instructions needed to reproduce the main experimental
523 results (either in the supplemental material or as a URL)? [No] Our experimental environment
524 are public and standard. All the information needed to reproduce our results is provided in the
525 main body and appendix. Code will be available publicly soon.
- 526 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)?
527 [Yes] Partially in main body and all details can be found in the appendix document.
- 528 (c) Did you report error bars (e.g., with respect to the random seed after running experiments
529 multiple times)? [Yes]
- 530 (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs,
531 internal cluster, or cloud provider)? [Yes]
- 532 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 533 (a) If your work uses existing assets, did you cite the creators? [Yes]
- 534 (b) Did you mention the license of the assets? [Yes] We use a free education licence for students for
535 MuJoCo.
- 536 (c) Did you include any new assets either in the supplemental material or as a URL? [No]
- 537 (d) Did you discuss whether and how consent was obtained from people whose data you’re us-
538 ing/curating? [N/A]
- 539 (e) Did you discuss whether the data you are using/curating contains personally identifiable informa-
540 tion or offensive content? [N/A]
- 541 5. If you used crowdsourcing or conducted research with human subjects...
- 542 (a) Did you include the full text of instructions given to participants and screenshots, if applicable?
543 [N/A]
- 544 (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB)
545 approvals, if applicable? [N/A]
- 546 (c) Did you include the estimated hourly wage paid to participants and the total amount spent on
547 participant compensation? [N/A]

548 **Appendix**

549 **A Supplementary Materials for Theoretical Analysis**

550 **A.1 More on Definition 1 and 2**

551 In Definition 1, we use \mathcal{P}_π for a formal description of value approximation process, i.e., the learning process
 552 of a parametrized PeVFA \mathbb{V}_θ with parameters $\theta \in \Theta$ to approximate the values of policy π . For a usual
 553 example, one can consider \mathcal{P}_π as multiple times of parameter update via gradient descent with respect to
 554 $f_\theta(\pi) = \|\mathbb{V}_\theta(\pi) - V^\pi\|$. Note that $f_\theta(\pi)$ can be equivalent to a common value approximation loss function
 555 $L(\theta) = \mathbb{E}_{s \sim \rho(s)} \left(\mathbb{V}_\theta(s, \pi) - \hat{V}^\pi(s) \right)^2$ with some unbiased estimates \hat{V}^π from experiences stored, when the
 556 same state distribution $\rho(s)$ is considered. Thus, with sufficient capacity of function approximation and certain
 557 number of training, we can expect a contraction of approximation loss for policy π obtained by \mathcal{P}_π .

558 We use Definition 2 to characterize the local smoothness of approximation loss f_θ near policy π with Lipschitz
 559 continuity. Consider a typical PeVFA V_θ parameterized by an MLP with finite weights, biases and non-linear
 560 activation. Such a V_θ is Lipschitz continuity with a bounded Lipschitz constant, as it is made up of function
 561 transformations that individually have bounded Lipschitz constants, e.g., weight matrix w of some layer has
 562 bounded Lipschitz constant to be the operator norm of matrix w and ReLU activation has Lipschitz constant of
 563 1. Further, easily we have for any π and π' ,

$$|f_\theta(\pi) - f_\theta(\pi')| \leq \|\mathbb{V}_\theta(\pi) - \mathbb{V}_\theta(\pi')\| + \|V^\pi - V^{\pi'}\|. \quad (2)$$

564 As mentioned above, V_θ is Lipschitz continuity with a bounded Lipschitz constant; and the norm of true value
 565 vector of two policies is also finite. Thus, f_θ in this case can also have a bounded Lipschitz constant L .

566 **A.2 Proof of Lemma 1**

567 *Proof.* For the clarity, we also use f and \hat{f} as abbreviations of f_θ and $f_{\hat{\theta}}$ in the following. Start from the
 568 \hat{L} -continuity of $\hat{f}(\theta)$ (recall Definition 2), we have the upper bound of $\hat{f}(\pi_2)$ below:

$$\hat{f}(\pi_2) \leq \hat{f}(\pi_1) + \hat{L} \cdot d(\pi, \pi'). \quad (3)$$

569 The second term in Equation 3 is decided by the two policies we considered and a Lipschitz constant \hat{L} . Moreover,
 570 the constant \hat{L} (i.e., locality property) is related to the parameters $\hat{\theta}$ of PeVFA. In general, we denote the above
 571 term as $\mathcal{M}(\pi_1, \pi_2, \hat{L})$ called *locality margin*. The locality margin $\mathcal{M}(\pi_1, \pi_2, \hat{L})$ can have different forms that
 572 depends on the specific locality property, for examples:

$$\mathcal{M}(\pi_1, \pi_2, \hat{L}) = \begin{cases} \hat{L} \cdot d(\pi_1, \pi_2) & \textcircled{1} \\ \langle \hat{f}'(\pi_1), \pi_2 - \pi_1 \rangle + \frac{1}{2} \hat{L} \cdot d(\pi_1, \pi_2)^2 & \textcircled{2} \\ \langle \hat{f}'(\pi_1), \pi_2 - \pi_1 \rangle + \frac{1}{2} \langle \hat{f}''(\pi_1)(\pi_2 - \pi_1), \pi_2 - \pi_1 \rangle + \frac{1}{6} \hat{L} \cdot d(\pi_1, \pi_2)^3 & \textcircled{3} \end{cases}$$

573 $\textcircled{1}$, $\textcircled{2}$, $\textcircled{3}$ correspond to Lipschitz Continuous, Lipschitz Gradients and Lipschitz Hessian [39], which are
 574 considered in previous works on generalization studies [22, 63].

575 Further, apply the Definition 1 and consider the case $f(\pi_1) \leq f(\pi_2)$, Equation 3 can be further transformed as
 576 follows:

$$\begin{aligned} \hat{f}(\pi_2) &\leq \hat{f}(\pi_1) + \mathcal{M}(\pi_1, \pi_2, \hat{L}) \\ &\leq \gamma f(\pi_1) + \mathcal{M}(\pi_1, \pi_2, \hat{L}) \\ &\leq \underbrace{\gamma f(\pi_2)}_{\text{generalized contraction}} + \underbrace{\mathcal{M}(\pi_1, \pi_2, \hat{L})}_{\text{locality margin}}, \end{aligned} \quad (4)$$

577 which yields the generalization upper bound in Lemma 1. We note the first term of RHS of Equation 4 as
 578 generalized contraction term since it is from the contraction on $f(\pi_1)$ caused by the value approximation
 579 operator \mathcal{P}_{π_1} , and the second term as locality margin since it is determined by specific local property. \square

580 **Remark 1** Since value approximation is only performed for π_1 , the condition $f_\theta(\pi_1) \leq f_\theta(\pi_2)$ can usually
 581 exist after a certain number of training; in turn, the complementary case $f_\theta(\pi_1) > f_\theta(\pi_2)$ is acceptable since
 582 the unoptimized approximation loss is already lower than the optimized one.

583 A.3 Proof of Corollary 1

584 *Proof.* Following Lemma 1, consider Lipschitz continuity for a concrete locality property of $f_{\hat{\theta}}$, we have,

$$\hat{f}(\pi_2) \leq \gamma f(\pi_2) + \hat{L} \cdot d(\pi_1, \pi_2). \quad (5)$$

585 Then we get the contraction condition of value generalization on π_2 in Corollary 1, by letting the RHS of
586 Equation 5 be smaller than $f(\pi_2)$:

$$\begin{aligned} \gamma f(\pi_2) + \hat{L} \cdot d(\pi_1, \pi_2) &< f(\pi_2) \\ (1 - \gamma)f(\pi_2) &> \hat{L} \cdot d(\pi_1, \pi_2) \\ f(\pi_2) &> \frac{\hat{L} \cdot d(\pi_1, \pi_2)}{1 - \gamma} \geq 0. \end{aligned} \quad (6)$$

587 □

588 **Remark 2** From the generalization contraction condition provided in Corollary 1, we can find that: as i. $\gamma \rightarrow 0$,
589 or ii. $d(\pi_1, \pi_2) \rightarrow 0$, or iii. $\hat{L} \rightarrow 0$, the contraction condition is easier to achieve (or the contraction gets
590 tighter), i.e., the generalization on unlearned policy π_2 is better.

591 In another word, the tighter the contraction on learned policy π_1 is, the closer the two policies are, the smoother
592 the approximation loss function \hat{f} is, the generalization on unlearned policy π_2 is better.

593 Corollary 1 provides the generalization contraction condition on $f(\pi_2)$, under the assumptions that \mathcal{P}_{π_1} is
594 γ -contraction and $f(\pi_1) < f(\pi_2)$ (as in Lemma 1). In below, we discuss a more general condition for
595 generalization contraction on $f(\pi_2)$ which indicates more possible cases:

596 **Corollary 3** For $\theta \xrightarrow{\mathcal{P}_{\pi_1}} \hat{\theta}$ and $f_{\hat{\theta}}$ is \hat{L} -continuous at π_1 , when $f(\pi_2) - \gamma f(\pi_1) > \hat{L} \cdot d(\pi_1, \pi_2)$, we have
597 that \mathcal{P}_{π_1} is also a γ_g -contraction for π_2 , i.e., $f_{\hat{\theta}}(\pi_2) \leq \gamma_g f_{\theta}(\pi_2)$ with $\gamma_g \in [0, 1)$.

598 *Proof.* From Equation 4, we have $\hat{f}(\pi_2) \leq \gamma f(\pi_1) + \hat{L} \cdot d(\pi_1, \pi_2)$. To yield the generalization contraction
599 on $f(\pi_2)$, is to let

$$\hat{f}(\pi_2) \leq \gamma f(\pi_1) + \hat{L} \cdot d(\pi_1, \pi_2) < f(\pi_2), \quad (7)$$

600 that is to let,

$$f(\pi_2) - \gamma f(\pi_1) > \hat{L} \cdot d(\pi_1, \pi_2). \quad (8)$$

601 □

602 Since $d(\pi_1, \pi_2)$ is constant in the two-policy case considered, the condition in Corollary 3 is associated to the
603 value approximation losses on π_1 and π_2 before applying the value approximation operator \mathcal{P}_{π} , as well as the
604 \hat{L} -continuity of $\hat{\theta}$ after applying \mathcal{P}_{π} . We can find similar conclusions as mentioned in Remark 2. However,
605 Corollary 3 indicates some more cases that the condition of generalization contraction can be satisfied. For
606 example, it can happen in the complementary cases as we assumed in Lemma 1, i.e., 1) when $f(\pi_1) > f(\pi_2)$,
607 or 2) \mathcal{P}_{π} is not a γ -contraction on $f(\pi_1)$.

608 A.4 Proof of Lemma 2

609 *Proof.* Consider any $t \geq 0$ and $\theta_{t-1} \xrightarrow{\mathcal{P}_{\pi_t}} \theta_t$, due to f_{θ_t} is L_t -continuous at π_t , we have,

$$f_{\theta_t}(\pi_{t+1}) \leq f_{\theta_t}(\pi_t) + L_t \cdot d(\pi_t, \pi_{t+1}), \quad (9)$$

610 then due to the definition of the value approximation process \mathcal{P}_{π_t} ,

$$\begin{aligned} f_{\theta_t}(\pi_{t+1}) &\leq \gamma_t f_{\theta_{t-1}}(\pi_t) + L_t \cdot d(\pi_t, \pi_{t+1}), \\ &= \gamma_t f_{\theta_{t-1}}(\pi_t) + \mathcal{M}_t, \end{aligned} \quad (10)$$

611 where $\mathcal{M}_t = L_t \cdot d(\pi_t, \pi_{t+1})$. □

612 Intuitively, such a recursive relation between the generalized approximation loss of two consecutive steps, i.e.,
613 $f_{\theta_{t-1}}(\pi_t)$ and $f_{\theta_t}(\pi_{t+1})$, are chained by the assumed continuity of the loss function f_{θ_t} and the definition of
614 value approximation process.

615 A.5 Proof of Corollary 2

616 *Proof.* Consider the consecutive value approximation process $\theta_{-1} \xrightarrow{\mathcal{P}^{\pi_0}} \theta_0 \xrightarrow{\mathcal{P}^{\pi_1}} \dots \xrightarrow{\mathcal{P}^{\pi_{t-1}}} \theta_{t-1} \xrightarrow{\mathcal{P}^{\pi_t}}$
 617 $\theta_t \xrightarrow{\mathcal{P}^{\pi_{t+1}}} \dots$, following the recursive relation in Lemma 2, we have the inequality below by induction,

$$\begin{aligned}
 f_{\theta_t}(\pi_{t+1}) &\leq \gamma_t f_{\theta_{t-1}}(\pi_t) + \mathcal{M}_t, \\
 &\leq \dots \\
 &\leq \gamma_t (\gamma_{t-1} (\dots (\gamma_0 f_{\theta_{-1}}(\pi_0) + \mathcal{M}_0) \dots) \mathcal{M}_{t-1}) + \mathcal{M}_t, \\
 &= \underbrace{\left(\prod_{i=0}^t \gamma_i \right) f_{\theta_{-1}}(\pi_0)}_{\mathbf{1}} + \underbrace{\sum_{i=0}^{t-1} \left(\prod_{j=i+1}^t \gamma_j \right) \mathcal{M}_i}_{\mathbf{2}} + \mathcal{M}_t.
 \end{aligned} \tag{11}$$

618 where $\mathcal{M}_t = L_t \cdot d(\pi_t, \pi_{t+1})$. We use $\mathbf{1}$ to denote the term for accumulated generalized contraction of initial
 619 approximation loss and use $\mathbf{2}$ to denote the term for accumulated locality margin. \square

620 Towards the infinity case i.e., $t \rightarrow \infty$, if we assume that (i) $\max_t d(\pi_t, \pi_{t+1}) < \infty$ and (ii) $\prod_{k=h_1}^{h_2} \gamma_k =$
 621 $O\left(\frac{1}{(h_2-h_1+1)^{1+\varepsilon}}\right)$, $\forall 0 < h_1 \leq h_2$ with some $\varepsilon > 0$, then $\lim_{t \rightarrow \infty} f_{\theta_t}(\pi_{t+1}) < \infty$. That is because the
 622 sequence $\{\mathcal{M}_i\}_{i=0}^t$ has a public upper bound $\mathcal{M}_{\max} = L_{\max} \cdot \max_t d(\pi_t, \pi_{t+1})$ where L_{\max} denotes the
 623 upper bound of Lipschitz constant (recall the discussion in Appendix A.1), and by (ii) $\sum_{i=0}^{t-1} \left(\prod_{j=i+1}^t \gamma_j \right) =$
 624 $O\left(\sum_{i=0}^{t-1} \frac{1}{(t-i+1)^{1+\varepsilon}}\right) < \infty$.

625 Note that we consider a really loose bound in the infinity case above with \mathcal{M}_{\max} , therefore the condition (ii) may
 626 be unnecessarily strict when the dynamics of L_t and $d(\pi_t, \pi_{t+1})$ are considered. Intuitively, the evolution of
 627 L_t during learning process is related to function family and optimization method of θ_t ; and for $d(\pi_t, \pi_{t+1})$, this
 628 is related to value approximation error ($f_{\theta_t}(\pi_t)$) and policy improvement method (i.e., how π_t is improved to be
 629 π_{t+1}). We leave these further analysis for future work.

630 A.6 Proof of Theorem 1

631 *Proof.* By the condition in Theorem 1, we have

$$\begin{aligned}
 f_{\theta_t}(\pi_t) + f_{\theta_t}(\pi_{t+1}) &\leq \|V^{\pi_t} - V^{\pi_{t+1}}\| \\
 &\leq \|\nabla_{\theta_t}(\pi_t) - V^{\pi_t}\| + \|\nabla_{\theta_t}(\pi_t) - V^{\pi_{t+1}}\| = f_{\theta_t}(\pi_t) + \|\nabla_{\theta_t}(\pi_t) - V^{\pi_{t+1}}\|,
 \end{aligned} \tag{12}$$

632 where the second inequality comes from *Triangle Inequality*. Then it is straightforward that

$$\underbrace{f_{\theta_t}(\pi_{t+1})}_{\text{generalized VAD with PeVFA}} \leq \underbrace{\|\nabla_{\theta_t}(\pi_t) - V^{\pi_{t+1}}\|}_{\text{conventional VAD}}, \tag{13}$$

633 which means that with local generalization of values for successive policy π_{t+1} , the value approximation distance
 634 (VAD) can be closer in contrast to the conventional one (RHS of Equation 13). \square

635 In practice, we consider that it is also possible for farther distance to exist, e.g., the condition in above Theorem 1
 636 is not satisfied. Moreover, under nonlinear function approximation, it is not necessary that a closer approximation
 637 distance (induced by Theorem 1) ensures easier approximation or optimization process. This can be associated
 638 to many factors, e.g., the underlying function space, the optimization landscape, the learning algorithm used
 639 and etc. In this paper, we provide a condition for potentially beneficial local generalization and we resort to
 640 empirical examination as shown in Sec. 3.3. Further investigation on the interplay between value generalization
 641 and policy learning especially under nonlinear function approximation is planned for future work.

642 B Details of Empirical Evidence of Two Kinds of Generalization

643 B.1 Global Generalization in 2D Point Walker

644 Global generalization denotes the generalization scenario that values can generalize to unlearned policies
 645 ($\pi' \in \Pi_1$) from already learned policies ($\pi \in \Pi_0$). We conduct the following experiments to demonstrate global
 646 generalization in a 2D continuous Point Walker environment with synthetic simple policies.

647 **Environment.** We consider a point walker on a 2D continuous plane with:

- 648 • state: $(x, y, \sin(\theta), \cos(\theta), \cos(x), \cos(y))$, where θ is the angle of the polar coordinates,
- 649 • action: 2D displacement, $a \in \mathbb{R}_{[-1,1]}^2$,

- a deterministic transition function that describes the locomotion of the point walker, depending on the current position and displacement issued by agent, i.e., $\langle x', y' \rangle = \langle x, y \rangle + a$,
- a reward function: $r_t = \frac{u_{t+1} - u_t}{10}$ with utility $u_t = x_t^2 - y_t^2$, as illustrated in Figure 7(a).

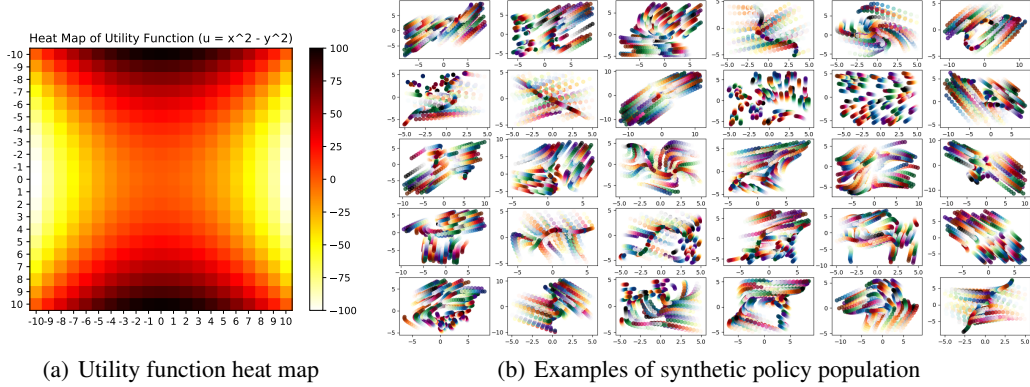


Figure 7: 2D Point Walker. (a) The heat map of the utility function of the 2D plane. The darker regions have higher utilities. (b) Demonstrative illustrations of trajectories generated by 30 synthetic policies, showing diverse behaviors and patterns. Each subplot illustrates the trajectories generated in 50 episodes by a randomly synthetic policy, with different colors as separation. For each trajectory (the same color in one subplot), transparency represents the dynamics along timesteps, i.e., fully transparent and non-transparent denotes the positions at first and last timesteps.

Synthetic Policy. We build the policy sets $\Pi = \Pi_0 \cup \Pi_1$ and $\Pi_0 \cap \Pi_1 = \emptyset$ with synthetic policies. Each synthetic policy is a 2-layer *tanh*-activated neural policy network with 2 nodes for each layer. The weights are initialized by sampling from a uniform distribution $U(-1, 1)$ and the biases are initialized by $U(-0.2, 0.2)$. Each policy is deterministic, taking an environmental state as input and outputting a displacement in the plane. We find that the synthetic population generated by such a simple way can show diverse behaviors. Figure 7(b) shows the motion patterns of an example of such a synthetic population. Note that the synthetic policies are not trained in this experiment.

Policy Dataset. We rollout each policy in environment to collect trajectories $\mathcal{T} = \{\tau_i\}_{i=0}^k$. For such small synthetic policies, it is convenient to obtain policy representation. Here we use the concatenation of all weights and biases of the policy network (26 in total) as representation χ_π for each policy π , called *raw* policy representation (RPR). Therefore, combined with the trajectories collected, we obtain the policy dataset, i.e., $\{(\chi_{\pi_j}, \mathcal{T}_{\pi_j})\}_{j=0}^n$. In total, 20k policies are synthesized in our experiments and we collected 50 trajectories with horizon 10 for each policy.

We separate the synthetic policies into training set (i.e., unknown policies Π_0) and testing set (i.e., unseen policies Π_1) in a proportion of 8 : 2. We set a PeVFA network $\mathbb{V}_\theta(s, \chi_\pi)$ to approximate the values of training policies (i.e., $\pi \in \Pi_0$), and then conduct evaluation on testing policies (i.e., $\pi \in \Pi_1$). We use Monte Carlo return [55] of collected trajectories as approximation target (true value of policies) in this experiment. The network architecture of $\mathbb{V}_\theta(s, \chi_\pi)$ is illustrated in Figure 8(a). The learning rate is 0.005, batch size is 256. K-fold validation is performed through shuffling training and testing sets.

Figure 8(b) shows the curves of training loss and testing loss. The average losses on training and testing set are 1.782 and 2.071 over 6 trials. Figure 2(a) plots the value predictions for policies from training and testing set (100 for each). This demonstrates that a PeVFA trained with data collected by training set Π_0 achieves reasonable value prediction of unseen testing policies in Π_1 . Our results indicate that value generalization can exist among policy space with a properly trained PeVFA. RPR can also be one alternative of policy representation when policy network is of small scale.

B.2 Local Generalization in MuJoCo Continuous Control Tasks

We demonstrate local generalization of PeVFA, especially to examine the existence of Theorem 1, i.e., PeVFA can induce closer approximation distance (i.e., lower approximation error) than conventional VFA along the policy improvement path.

We use a 2-layer 8-unit policy network trained by PPO [50] algorithm in OpenAI MuJoCo continuous control tasks. As in previous section, using a very small policy network is for the convenience of training and acquisition

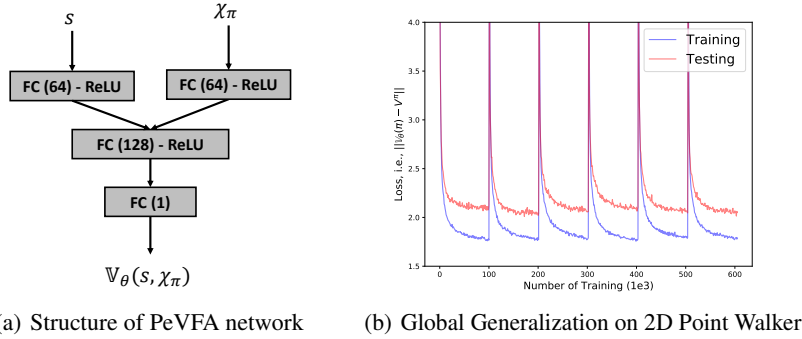


Figure 8: Global generalization of PeVFA on 2D Point Walker. (a) An illustration of architecture of PeVFA network. FC is abbreviation for Fully-connected layer. (b) Training and testing losses. Data shuffle and network re-initialization are performed per 100 steps, i.e., 1e5 training times.

684 of policy representation in this demonstrative experiment. We use all weights and biases of the small policy
 685 network (also called *raw* policy representation, RPR), whose number is about 10 to 100 in our experiments,
 686 depending on the specific environment (i.e., the state and action dimensions). We train the small policy network
 687 as commonly done with PPO [50] and GAE [49]. The conventional value network $V_\phi(s)$ (VFA), is a 2-layer
 688 128-unit ReLU-activated MLP with state as input and value as output. Parallel to the conventional VFA in PPO,
 689 we set a PeVFA network $V_\theta(s, \chi_\pi)$ with RPR as additional input. The structure of PeVFA differs at the first
 690 hidden layer which has two input streams and each of them has 64 units, as illustrated in Figure 8(a), so that
 691 making VFA and PeVFA have similar scales of parameter number. In contrast to conventional VFA V_ϕ which
 692 approximates the value of current policy (e.g., Algorithm 2), PeVFA $V_\theta(s, \chi_\pi)$ has the capacity to preserve
 693 values of multiple policies and thus is additionally trained to approximate the values of all historical policies
 694 $(\{\pi_i\}_{i=0}^t)$ along the policy improvement path (e.g., Algorithm 3). The learning rate of policy is 0.0001 and
 695 the learning rate of value function approximators ($V_\phi(s)$ and $V_\theta(s, \chi_\pi)$) is 0.001. The training scheme of PPO
 696 policy here is the same as that described in Appendix F.1 and Table 2.

697 Note that $V_\theta(\chi_\pi)$ does not interfere with PPO training here, and is only referred as a comparison with V_ϕ on
 698 the approximation error to the true values of successive policy π_{t+1} . We use the MC returns of on-policy data
 699 (i.e., trajectories) collected by current successive policy as unbiased estimates of true values, similarly done
 700 in [61, 12]. Then we calculate the approximation error for VFA V_ϕ and PeVFA $V_\theta(\chi_\pi)$ to the approximation
 701 target before and after value network training of current iteration. Finally, we compare the approximation error
 702 between VFA and PeVFA to approximately examine local generalization and closer approximation target in
 703 Theorem 1. Complete results of local generalization across all 7 MuJoCo tasks are show in Figure 9. The
 704 results show that PeVFA consistently shows lower losses (i.e., closer to approximation target) across all tasks
 705 than conventional VFA before and after policy evaluation along policy improvement path, which demonstrates
 706 Theorem 1. Moreover, we also provide similar empirical evidence when policy is updated with larger learning
 707 rates in $\{0.0001, 0.001, 0.005\}$, as in Figure 10.

708 A common observation across almost all results in Figure 9 and in Figure 10 is that the larger the extent of policy
 709 change (see the regions with a sheer slope on green curves), the higher the losses of conventional VFA tend
 710 to be (see the peaks of red curves), where the generalization tends to be better and more significant (see the
 711 blue curves). Since InvertedPendulum-v1 is a simple task while the complexity of the solution for Ant-v1 is
 712 higher, the difference between value approximation losses of PeVFA and VFA is more significant at the regions
 713 with fast policy improvement. Besides, the Raw Policy Representation (RPR) we used here does not necessarily
 714 induce a smooth and efficient policy representation space, among which policy values are easy to generalize
 715 and optimize. Thus, RPR may be sufficient for a good generalization in InvertedPendulum-v1 but may be not
 716 in Ant-v1. Overall, we think that the quantity of value approximation loss is related to several factors of the
 717 environment such as the reward scale, the extent of policy change, the complexity of underlying solution (e.g.,
 718 value function space) and some others. A further investigation on this can be interesting.

719 C Generalized Policy Iteration with PeVFA

720 C.1 Comparison between Conventional GPI and GPI with PeVFA

721 A graphical comparison of conventional GPI and GPI with PeVFA is shown in Figure 1. Here we provide
 722 another comparison with pseudo-codes.

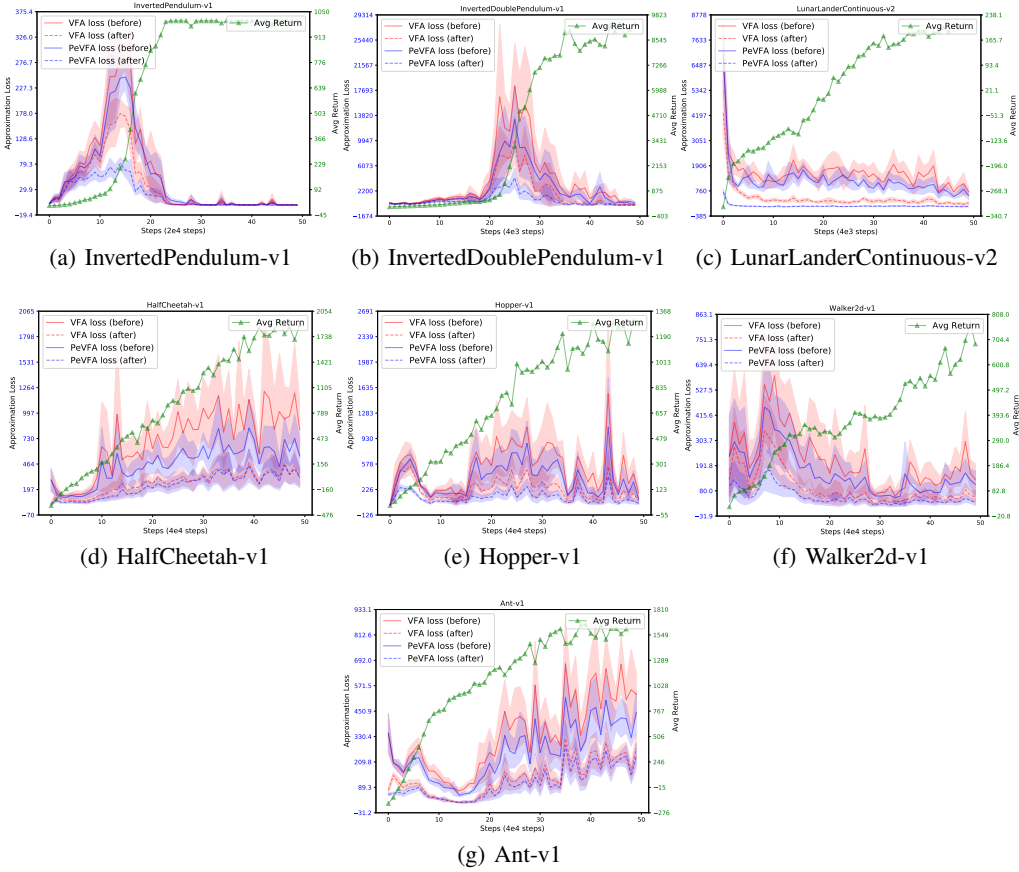


Figure 9: Complete empirical evidence of local generalization of PeVFA across 7 MuJoCo tasks. The learning rate of policy and value function approximators are 0.0001 and 0.001 respectively. Each plot has two vertical axes, the left one for approximation error (red and blue curves) and the right one for average return (green curves). Red and blue denotes the approximation error of conventional VFA ($V_\phi(s)$) and of PeVFA ($\nabla_\theta(s, \chi_\pi)$) respectively; solid and dashed curves denote the approximation error before and after the training for values of successive policy (i.e., policy evaluation) with conventional VFA and PeVFA, averaged over 6 trials. The shaded region denotes half a standard deviation of average evaluation. PeVFA consistently shows lower losses (i.e., closer to approximation target) across all tasks than conventional VFA before and after policy improvement along policy improvement path, which demonstrates Theorem 1.

723 From the lens of Generalized Policy Iteration [55], for most model-free policy-based RL algorithms, the ap-
 724 proximation of value function and the update of policy through policy gradient theorem are usually conducted
 725 iteratively. Representative examples are REINFORCE [55], Advantage Actor-Critic [36], Deterministic Policy
 726 Gradient (DPG) [52] and Proximal Policy Optimization (PPO) [50]. With conventional value function (approx-
 727 imator), policy evaluation is usually performed in an on-policy or off-policy fashion. We provide a general GPI
 728 description of model-free policy-based RL algorithm with conventional value functions in Algorithm 2.

729 Note that we use subscript $t-1 \rightarrow t$ (Line 13 in Algorithm 2) to let the updated value functions to correspond
 730 to the evaluated policy π_t during policy evaluation process in current iteration.

731 As a comparison, a new form of GPI with PeVFA is shown in Algorithm 3. Except for the different parameteriza-
 732 tion of value function, PeVFA can perform additionally training on historical policy experiences at each iteration
 733 (Line 7-8). This is naturally compatible with PeVFA since it develops the capacity of conventional value function
 734 to preserve the values of multiple policies. Such a training is to improve the value generalization of PeVFA
 735 among a policy set or policy space. Note that for value approximation of current policy π_t (Line 10-14), the
 736 start points are generalized values of π_t from historical approximation, i.e., $\nabla_{t-1}(s, \chi_{\pi_t})$ and $Q_{t-1}(s, a, \chi_{\pi_t})$.
 737 In another word, this is the place where local generalization steps (illustrated in Figure 2(b)) are. One may
 738 compare with conventional start points ($V_{t-1}^\pi(s)$ and $Q_{t-1}^\pi(s, a)$, Line 13 in Algorithm 2) and see the difference,

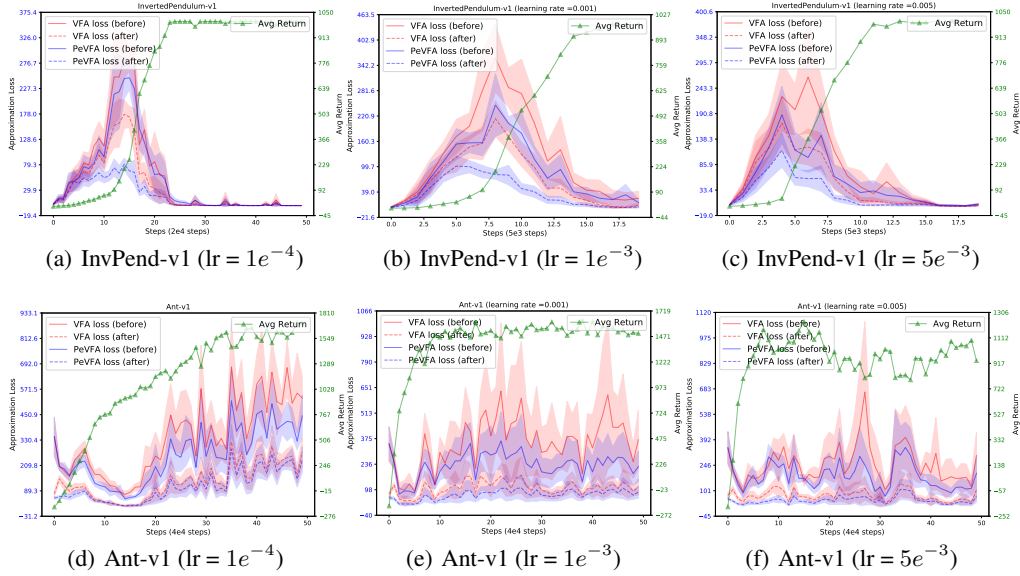


Figure 10: Empirical evidence of local generalization of PeVFA on InvertedPendulum-v1 and Ant-v1 with different learning rates of policy, i.e., $\{0.0001, 0.001, 0.005\}$. Results are averaged over 6 trials.

Algorithm 2 Generalized policy iteration for model-free policy-based RL algorithm with conventional value functions ($V^\pi(s)$ or $Q^\pi(s, a)$)

- 1: Initialize policy π_0 and $V_{-1}^\pi(s)$ or $Q_{-1}^\pi(s, a)$
- 2: Initialize experience buffer \mathcal{D}
- 3: **for** iteration $t = 0, 1, 2, \dots$ **do**
- 4: Rollout policy π_t in the environment and obtain trajectories $\mathcal{T}_t = \{\tau_i\}_{i=0}^k$
- 5: Add experiences \mathcal{T}_t in buffer \mathcal{D}
- 6: **if on-policy update then**
- 7: Prepare training samples from rollout trajectories \mathcal{T}_t
- 8: **else if off-policy update then**
- 9: Prepare training samples by sampling from buffer \mathcal{D}
- 10: **end if**
- 11: Calculate approximation target $\{y_i\}_i$ from training samples (e.g., with MC or TD)
- 12: # Generalized Policy Evaluation
- 13: Update $V_{t-1}^\pi(s)$ or $Q_{t-1}^\pi(s, a)$ with $\{(s_i, y_i)\}_i$ or $\{(s_i, a_i, y_i)\}_i$, i.e., $V_t^\pi \leftarrow V_{t-1}^\pi$ or $Q_t^\pi \leftarrow Q_{t-1}^\pi$
- 14: # Generalized Policy Improvement
- 15: Update policy π_t with regard to $V_t^\pi(s)$ or $Q_t^\pi(s, a)$ through some policy gradient theorem, i.e., $\pi_{t+1} \leftarrow \pi_t$
- 16: **end for**

739 e.g., $V_{t-1}^\pi(s) \Leftrightarrow V^{t-1}(s) \Leftrightarrow \mathbb{V}_{t-1}(s, \chi_{\pi_{t-1}})$ is different with $\mathbb{V}_{t-1}(s, \chi_{\pi_t})$, where \Leftrightarrow is used to denote an
740 equivalence in definition. As discussed in Sec. 3.3 and 3.4, we suggest that such local generalization steps help
741 to reduce approximation error and thus improve efficiency during the learning process.

742 C.2 More Discussions on GPI with PeVFA

743 **Off-Policy Learning.** Off-policy Value Estimation [55] denotes to evaluate the values of some target policy
744 from data collected by some behave policy. As commonly seen in RL (also shown in Line 6-10 in Algorithm
745 2), different algorithms adopt on-policy or off-policy methods. For GPI with PeVFA, especially for the value
746 estimation of historical policies (Line 8 in Algorithm 3), on-policy and off-policy methods can also be considered
747 here. One interesting thing is, in off-policy case, one can use experiences from any policy for the learning
748 of another one, which can be appealing since the high data efficiency of value estimation of each policy can

Algorithm 3 Generalized policy iteration of model-free policy-based RL algorithm with PeVFAs ($\mathbb{V}(s, \chi_\pi)$ or $\mathbb{Q}(s, a, \chi_\pi)$)

```

1: Initialize policy  $\pi_0$  and PeVFA  $\mathbb{V}_{-1}(s, \chi_\pi)$  or  $\mathbb{Q}_{-1}(s, a, \chi_\pi)$ 
2: Initialize experience buffer  $\mathcal{D}$ 
3: for iteration  $t = 0, 1, 2, \dots$  do
4:   Rollout policy  $\pi_t$  in the environment and obtain trajectories  $\mathcal{T}_t = \{\tau_i\}_{i=0}^k$ 
5:   Get the policy representation  $\chi_{\pi_t}$  for policy  $\pi_t$  (from policy network parameters or policy rollout experiences)
6:   Add experiences  $(\chi_{\pi_t}, \mathcal{T}_t)$  in buffer  $\mathcal{D}$ 
7:   # Value approximation training for historical policies  $\{\pi_i\}_{i=0}^{t-1}$ 
8:   Update PeVFA  $\mathbb{V}_{t-1}(s, \chi_{\pi_i})$  or  $\mathbb{Q}_{t-1}(s, a, \chi_{\pi_i})$  with all historical policy experiences  $\{(\chi_{\pi_i}, \mathcal{T}_i)\}_{i=0}^{t-1}$ 
9:   # Conventional value approximation training for current policy  $\pi_t$ 
10:  if on-policy update then
11:    Update PeVFA  $\mathbb{V}_{t-1}(s, \chi_{\pi_t})$  or  $\mathbb{Q}_{t-1}(s, a, \chi_{\pi_t})$  for  $\pi_t$  with on-policy experiences  $(\chi_{\pi_t}, \mathcal{T}_t)$ 
12:  else if off-policy update then
13:    Update PeVFA  $\mathbb{V}_{t-1}(s, \chi_{\pi_t})$  or  $\mathbb{Q}_{t-1}(s, a, \chi_{\pi_t})$  for  $\pi_t$  with off-policy experiences  $\chi_{\pi_t}$  and  $\{\mathcal{T}_i\}_{i=0}^t$  from experience buffer  $\mathcal{D}$ 
14:  end if
15:   $\mathbb{V}_t \leftarrow \mathbb{V}_{t-1}$  or  $\mathbb{Q}_t \leftarrow \mathbb{Q}_{t-1}$ 
16:  Update policy  $\pi_t$  with regard to  $\mathbb{V}_t(s, \chi_{\pi_t})$  or  $\mathbb{Q}_t(s, a, \chi_{\pi_t})$  through some policy gradient theorem, i.e.,  $\pi_{t+1} \leftarrow \pi_t$ 
17: end for

```

749 strengthen value generalization among themselves with PeVFA, which further improve the value estimation
750 process.

751 **Convergence of GPI with PeVFA.** Convergence of GPI is usually discussed in some ideal cases, e.g., with
752 small and finite state action spaces and with sufficient function approximation ability. In this paper, we focus on
753 the comparison between conventional VFA and PeVFA in value estimation, i.e., Policy Evaluation, and we make
754 no assumption on the Policy Improvement part. We conjecture that with the same policy improvement algorithm
755 and sufficient function approximation ability, GPI with conventional VFA and GPI with PeVFA finally converge
756 to the same policy. Moreover, based on Theorem 1 and our empirical evidence in Sec. 3.3, GPI with PeVFA can
757 be more efficient in some cases: with local generalization, it could take less experiences (training) for PeVFA to
758 reach the same level of approximation error than conventional VFA, or with the same amount of experience
759 (training), PeVFA could achieve lower approximation error than conventional VFA. We believe that a deeper
760 dive in convergence analysis is worth further investigation.

761 **PeVFA with TD Value Estimation.** In this paper, we propose PPO-PeVFA as a representative instance of
762 re-implementing DRL algorithms with PeVFA. Our theoretical results and algorithm 3 proposed under the
763 general policy iteration (GPI) paradigm are suitable for TD value estimation as well in principle. One potential
764 thing that deserves further investigation is that, it can be a more complex generalization problem since the
765 approximation target of TD learning is moving (in contrast to the stationary target when unbiased Monte Carlo
766 estimates are used). The non-stationarity induced by TD is recognized to hamper the generalization performance
767 in RL as pointed out in recent work [21]. Further study on PeVFA with TD learning (e.g., TD3 and SAC) is
768 planned in the future as mentioned in Sec. 6.

769 D Policy Representation Learning Details

770 D.1 Policy Geometry

771 A policy $\pi \in \Pi = \mathcal{P}(\mathcal{A})^{\mathcal{S}}$, defines the behavior (action distribution) of the agent under each state. For a more
772 intuitive view, we consider the geometrical shape of a policy: all state $s \in \mathcal{S}$ and all action $a \in \mathcal{A}$ are arranged
773 along the x -axis and y -axis of a 2-dimensional plane, and the probability (density) $\pi(a|s)$ is the value of z -axis
774 over the 2-dimensional plane. Note that for finite state space and finite action space (discrete action space), the
775 policy can be viewed as a $|\mathcal{S}| \times |\mathcal{A}|$ table with each entry in it is the probability of the corresponding state-action
776 case. Without loss of generality, we consider the continuous state and action space and the policy geometry here.
777 Illustrations of policy geometry examples are shown in Figure 11.

778 Figure 11(a) shows the policy geometry in a general case, where the policy can be defined arbitrarily. Generally,
 779 the policy geometry can be any possible geometrical shape (s.t. $\forall s \in \mathcal{S}, \sum_{a \in \mathcal{A}} \pi(a|s) = 1$). This means that
 780 the policy geometry is not necessarily continuous or differentiable in a general case. Specially, one can imagine
 781 that the geometry of a deterministic policy consists of peak points ($z = 1$) for each state and other flat regions
 782 ($z = 0$). Figure 11(b) shows an example of synthetic continuous policy which can be viewed as a 3D curved
 783 surface. In Deep RL, a policy may usually be modeled as a deep neural network. Assume that the neural policy
 784 is a function that is almost continuous and differentiable everywhere, the geometry of such a neural policy can
 785 also be continuous and differentiable almost everywhere. As shown in Figure 11(c), we provide a demo of neural
 786 policy by smoothing an arbitrary policy along both state and action axes.




Figure 11: Examples of policy geometry. (a) An arbitrary policy, where $p(s, a)$ is sampled from $\mathcal{N}(0, 1)$ for a joint space of 40 states and 40 actions and then normalized along action axis. States are squeezed into the range of $[-1, 1]$ for clarity. (b) A synthetic continuous policy with $p(s, a) = (1 - a^5 + s^5) \exp(-s^2 - a^2)$ for a joint space of $s \in [-2, 2]$ and $a \in [-2, 2]$ (each of which are discretized into 40 ones) and then normalized along action axis. (c) A general demo of neural network policy, generated from an arbitrary policy (as in (a)) over a joint space of 200 states and 100 actions with some smoothing skill. States are squeezed into the range of $[-1, 1]$ for clarity and the probability masses of actions under each state are normalized to sum into 1.

787 D.2 Implementation Details of Surface Policy Representation (SPR) and Origin Policy 788 Representation (OPR)

789 Here we provide a detailed description of how to encode different policy data for Surface Policy Representation (SPR) and Origin Policy Representation (OPR) we introduced in Sec. 4.

791 **Encoding of State-action Pairs for SPR.** Given a set of state-action pairs $\{s_i, a_i\}_{i=1}^n$ (with size $[n, s_dim + a_dim]$) generated by policy π (i.e., $a_i \sim \pi(\cdot|s_i)$), we concatenate each state-action pair and obtain an embedding of it by feeding it into an MLP, resulting in a stack of state-action embedding with size $[n, e_dim]$.
 793 After this, we perform a mean-reduce operator on the stack and obtain an SPR with size $[1, e_dim]$. A similar
 794 permutation-invariant transformation is previously adopted to encode trajectories in [14].
 795

796 **Encoding of Network Parameters for OPR.** We propose a novel way to learn low-dimensional embedding
 797 from policy network parameters directly. To our knowledge, we are the first to learn policy embedding from
 798 neural network parameters in RL. Note that latent space of neural networks are also studied in differentiable
 799 Network Architecture Search (NAS) [32, 33], where architecture-level embedding are usually considered. In
 800 contrast, OPR cares about parameter-level embedding with a given architecture.

801 Consider a policy network to be an MLP with well-represented state (e.g., CNN for pixels, LSTM for sequences)
 802 as input and deterministic or stochastic policy output. We compress all the weights and biases of the MLP to
 803 obtain an OPR that represents the decision function. The encoding process of an MLP with two hidden layers
 804 is illustrated in Figure 12. The main idea is to perform permutation-invariant transformation for inner-layer
 805 weights and biases for each layer first. For each unit of some layer, we view the unit as a non-linear function of
 806 all outputs, determined by weights, a bias term and activation function. Thus, the whole layer can be viewed
 807 as a batch of operations of previous outputs, e.g., with the shape $[h_t, h_{t-1} + 1]$ for $t \geq 1$ and $t = 0$ is also for
 808 the input layer. Note that we neglect activation function in the encoding since we consider the policy network
 809 structure is given. That is also why we call OPR as parameter-level embedding in contrast to architecture-level
 810 embedding in NAS (mentioned in the last paragraph). We then feed the operation batch in an MLP and perform
 811 mean-reduce to outputs. Finally we concatenate encoding of layers and obtain the OPR.

812 We use permutation-invariant transformation for OPR because that we suggest the operation batch introduced
 813 in the last paragraph can be permutation-invariant. Actually, our encoding shown in Figure 12 is not strict to
 814 obtain permutation-invariant representation since inter-layer dependencies are not included during the encoding

815 process. We also tried to incorporate the order information during OPR encoding and we found similar results
 816 with the way we present in Figure 12, which we adopt in our experiments.

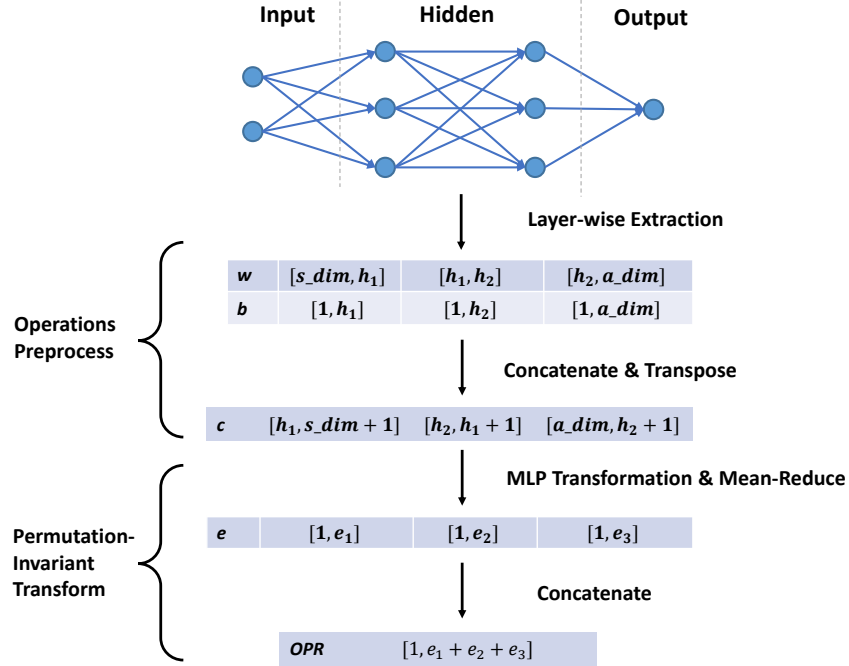


Figure 12: An illustration for policy encoder of Origin Policy Representation (OPR) for a two-layer MLP. h_1, h_2 denotes the numbers of hidden units for the first and second hidden layers respectively. The main idea is to perform permutation-invariant transformation for inner-layer weights and biases for each layer first and then concatenate encoding of layers.

817 **Towards more sophisticated RL policy that operates images.** Our proposed two policy representations (i.e.,
 818 OPR and SPR) can basically be applied to encode policies that operate images, with the support of advanced
 819 image-based state representation. For OPR, a policy network with image input usually has a pixel feature
 820 extractor like Convolutional Neural Networks (CNNs) followed by a decision model (e.g., an MLP). With
 821 effective features extracted, the decision model can be of moderate (or relatively small) scale. Recent works on
 822 unsupervised representation learning like MoCo [20], SimCLR [7], CURL [54] also show that a linear classifier
 823 or a simple MLP which takes compact representation of images learned in an unsupervised fashion is capable of
 824 solving image classification and image-based continuous control tasks. In another direction, it is promising to
 825 develop more efficient even gradient-free OPR, for example using the statistics of network parameters in some
 826 way instead of all parameters as similarly considered in [60].

827 For SPR, to encode state-action pairs (or sequences) with image states can be converted to the encoding in
 828 the latent space. The construction of latent space usually involves self-supervised representation learning, e.g.,
 829 image reconstruction, dynamics prediction. A similar scenario can be found in recent model-based RL like
 830 Dreamer [16], where the imagination is efficiently carried out in the latent state space rather than among original
 831 image observations.

832 Overall, we believe that there remain more effective approaches to represent RL policy to be developed in the
 833 future in a general direction of OPR and SPR, which are expected to induce better value generalization in a
 834 different RL problems.

835 D.3 Data Augmentation for SPR and OPR in Contrastive Learning

836 Data augmentation is studied to be an important component in contrastive learning in deep RL recently [24, 28].
 837 Contrastive learning usually resorts to data augmentation to build positive samples. Data augmentation is
 838 typically performed on pixel inputs (e.g., images) problems [20, 7]. In our work, we train policy representation
 839 with contrastive learning where data augmentation is performed on policy data. For SPR, i.e., state-action pairs
 840 as policy data, there is no need to perform data augmentation since different batches of randomly sampled
 841 state-action pairs naturally forms positive samples, since they all reflect the behavior of the same policy. A
 842 similar idea can also be found in [11] when dealing with task context in Meta-RL.

843 For OPR, i.e., policy network parameters as policy data, it is unclear how to perform data augmentation on them.
 844 In this work, we consider two kinds of data augmentation for policy network parameters as shown in Figure 13.
 845 We found similar results for both random mask and noise corruption, and we use random mask as default data
 846 augmentation in our experiments.




Figure 13: Examples of data augmentation on policy network parameters for Origin Policy Representation (OPR). *Left:* an example of original policy network. *Middle:* dropout-like random masks are performed on original policy network, where gray dashed lines represent the weights masked out. *Right:* randomly selected weights are corrupted by random noises, denoted by orange lines.

847 As an unsupervised representation learning method, contrastive Learning encourages policies to be close to
 848 similar ones (i.e., positive samples π^+) and to be apart from different ones (i.e., negative samples π^-) in policy
 849 representation space. The policy representation network is then trained with InfoNCE loss [41], i.e., to minimize
 850 the cross-entropy loss below:

$$\mathcal{L}_{CL} = -\mathbb{E} \left[\log \frac{\exp(\chi_{\pi}^T W \chi_{\pi^+})}{\exp(\chi_{\pi}^T W \chi_{\pi^+}) + \sum_{\pi^-} \exp(\chi_{\pi}^T W \chi_{\pi^-})} \right]$$

851 D.4 Pseudo-code of Policy Representation Learning Framework

852 The pseudo-code of the overall framework of policy representation learning is in Algorithm 4. The policy
 853 representation learning is conducted base on a policy dataset, which stores the policy data, i.e., interaction
 854 trajectories generated by policies and the parameters of policy networks. Such a dataset can be obtained in
 855 different ways, e.g., pre-collected, online collected and etc. In our experiments, we do not assume the access
 856 to pre-collected or given policy data; instead, we use the data of all historical policies met along the policy
 857 improvement path during the online learning process.

858 Different kinds of policy data (i.e., state-action pairs or policy parameters) are used depending on the policy
 859 representation adopted (i.e., SPR or OPR). For policy representation learning, the value function approximation
 860 loss (E2E) is used as a default choice of training loss in our framework. In addition, the auxiliary loss (AUX) of
 861 policy recovery and contrastive learning (CL) serve as another two options to be optimized for representation
 862 learning. Note that in Line 21, the positive samples $\chi_{\pi_i^+}$ is obtained from a momentum policy encoder [20] with
 863 another augmentation for corresponding policy data, while negative samples $\chi_{\pi_i^-}$ are other policy embeddings
 864 in the same batch, i.e., $\chi_{\pi_i^-} \in B \setminus \{\chi_{\pi_i}\}$.

865 D.5 Criteria of A Good Policy Representation

866 To answer the question: what is a good representation for RL policy ought to be? We assume the following
 867 criteria:

- 868 • *Dynamics.* Intuitively, a good policy representation should contain the information of how the policy
 869 influences the environment (dynamics and rewards).
- 870 • *Consistency.* A good policy representation should keep the consistency among both policy space and
 871 presentation space. Concretely, the policy representation should be distinguishable, i.e., different
 872 policies also differ among their representation. In contrast, the representation of similar polices should
 873 lay on the close place in the representation space.
- 874 • *Geometry.* Additionally, from the lens of policy geometry as shown in Appendix D.1, a good policy
 875 representation should be an reflection of policy geometry. It should show a connection to the policy
 876 geometry or be interpretable from the geometric view.

877 From the perspective of above criteria, SPR follows *Dynamics* and *Geometry* while OPR may render them in
 878 an implicit way since network parameters determine the nonlinear function of policy. Auxiliary loss for policy

Algorithm 4 A Framework of Policy Representation Learning

Input: policy dataset $\mathbb{D} = \{(\pi_i, \omega_i, \mathcal{D}_{\pi_i})\}_{i=1}^n$, consisting of policy π_i , policy parameters ω_i and state-action pairs $\mathcal{D}_{\pi_i} = \{(s_j, a_j)\}_{j=1}^m$

- 1: Initialize the policy encoder g_α with parameters α
 - 2: Initialize the policy decoder (or master policy) (network) $\bar{\pi}_\beta(a|s, \chi_\pi)$ for SPR and the weight matrix W for ORP respectively
 - 3: **for** iteration $i = 0, 1, 2, \dots$ **do**
 - 4: Sample a mini-batch of policy data \mathcal{B} from \mathbb{D}
 - 5: # Encode and obtain the policy embedding χ_{π_i} with SPR or OPR
 - 6: **if** Use OPR **then**
 - 7: **if** Use Contrastive Learning **then**
 - 8: Perform data augmentation on each $w_i \in \mathcal{B}$
 - 9: **end if**
 - 10: $\chi_{\pi_i} = g_\alpha^{\text{OPR}}(\omega_{\pi_i})$ for each $(\pi_i, \omega_i, \cdot) \in \mathcal{B}$
 - 11: **else if** Use SPR **then**
 - 12: $\chi_{\pi_i} = g_\alpha^{\text{SPR}}(B_i)$ where B_i is a mini-batch of state-action pairs sampled from \mathcal{D}_{π_i} , for each $(\pi_i, \cdot, \mathcal{D}_{\pi_i}) \in \mathcal{B}$
 - 13: **end if**
 - 14: # Train policy encoder g_α in different ways (i.e., AUX or CL)
 - 15: **if** Use Auxiliary Loss (AUX) **then**
 - 16: Sample a mini-batch of state-action pairs $B = (s_i, a_i)_{i=1}^b$ from \mathcal{D}_{π_i} for each π_i
 - 17: Compute the auxiliary loss, $\mathcal{L}_{\text{Aux}} = -\sum_{(s_i, a_i) \in B} \log \bar{\pi}_\alpha(a_i|s_i, \chi_{\pi_i})$
 - 18: Update parameters α, β to minimize \mathcal{L}^{Aux}
 - 19: **end if**
 - 20: **if** Use Contrastive Learning (CL) **then**
 - 21: Calculate contrastive loss, $\mathcal{L}_{\text{CL}} = -\sum_{\chi_{\pi_i} \in B} \log \frac{\exp(\chi_{\pi_i}^T W \chi_{\pi_i^+})}{\exp(\chi_{\pi_i}^T W \chi_{\pi_i^+}) + \sum_{\pi_i^-} \exp(\chi_{\pi_i}^T W \chi_{\pi_i^-})}$,
 - where $\chi_{\pi_i^+}, \chi_{\pi_i^-}$ are positive and negative samples
 - 22: Update parameters α, W to minimize \mathcal{L}^{CL}
 - 23: **end if**
 - 24: # Train policy encoder g_α with the PeVFA approximation loss (E2E)
 - 25: Calculate the value approximation loss of PeVFA, \mathcal{L}_{Val}
 - 26: Update parameters α to minimize \mathcal{L}_{Val}
 - 27: **end for**
-

879 recovery (AUX) is a learning objective to acquire *Dynamics*; Contrastive Learning (CL) is used to impose
880 *Consistency*.

881 Based on the above thoughts, we hypothesize the reasons of several findings as shown in the comparison
882 in Table 1. First, AUX naturally overlaps with SPR and OPR to some degree for *Dynamics* while CL is
883 relatively complementary to SPR and OPR for *Consistency*. This may be the reason why CL improves the E2E
884 representation more than AUX in an overall view. Second, the noise of state-action samples for SPR may be
885 the reason to OPR’s slightly better overall performance than that of SPR (similar results are also found in our
886 visualizations as in Figure 19).

887 Moreover, the above criteria are mainly considered from an unsupervised or self-supervised perspective. However,
888 a sufficiently good representation of all the above properties may not be necessary for a specific downstream
889 generalization or learning problem which utilizes the policy representation. A problem-specific learning signal,
890 e.g., the value approximation loss in our paper (E2E representation), can be efficient since it is to extract the
891 most relevant information in policy representation for the problem. A recent work [59] also studies the relation
892 between self-supervised representation and downstream tasks from the lens of mutual information. Therefore,
893 we suggest that a trade-off between good unsupervised properties and efficient problem-specific information of
894 policy representation should be considered when using policy representation in a specific problem.

895 E Complete Background and Detailed Related Work

896 E.1 Reinforcement Learning

897 **Markov Decision Process.** We consider a Markov Decision Process (MDP) defined as $\langle \mathcal{S}, \mathcal{A}, r, \mathcal{P}, \gamma, \rho_0 \rangle$
898 with \mathcal{S} the state space, \mathcal{A} the action space, r the reward function, \mathcal{P} the transition function, $\gamma \in [0, 1)$ the
899 discount factor and ρ_0 the initial state distribution. A policy $\pi \in \Pi = P(\mathcal{A})^{|\mathcal{S}|}$, defines the distribution over
900 all actions for each state. The agent interacts with the environment with its policy, generating the trajectory
901 $s_0, a_0, r_1, s_1, a_1, r_2, \dots, s_t, a_t, r_{t+1}, \dots$, where $r_{t+1} = r(s_t, a_t)$. An RL agent seeks for an optimal policy that
902 maximizes the expected long-term discounted return, $J(\pi) = \mathbb{E}_{s_0 \sim \rho_0, a \sim \pi} [\sum_{t=0}^{\infty} \gamma^t r_{t+1}]$.

903 **Value Function.** Almost all RL algorithms involve value functions [55], which estimate how good a state or
904 a state-action pair is conditioned on a given policy. The *state-value function* $v^\pi(s)$ is defined in terms of the
905 expected return obtained through following the policy π from a state s :

$$v^\pi(s) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r_{t+1} | s_0 = s \right] \text{ for all } s \in \mathcal{S}.$$

906 Similarly, *action-value function* is defined for all state-action pairs as $q^\pi(s, a) =$
907 $\mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r_{t+1} | s_0 = s, a_0 = a]$. Typically, value functions are learned through Monte Carlo (MC)
908 or Temporal Difference (TD) algorithms [55].

909 Bellman equations defines the recursive relationships among value functions. The *Bellman Expectation equation*
910 of $v^\pi(s)$ has a matrix form as below [55]:

$$V^\pi = r^\pi + \gamma \mathcal{P}^\pi V^\pi = (\mathcal{I} - \gamma \mathcal{P}^\pi)^{-1} r^\pi, \quad (14)$$

911 where V^π is a $|\mathcal{S}|$ -dimensional vector, \mathcal{P}^π is the state-to-state transition matrix $\mathcal{P}^\pi(s'|s) =$
912 $\sum_{a \in \mathcal{A}} \pi(a|s) \mathcal{P}(s'|s, a)$ and r^π is the vector of expected rewards $r^\pi(s) = \sum_{a \in \mathcal{A}} \pi(a|s) r(s, a)$. Equation 14
913 indicates that value function is determined by policy π and environment models (i.e., \mathcal{P} and r). For a conventional
914 value function, all of them are modeled implicitly within a table or a function approximator, i.e., a mapping from
915 only states (and actions) to values.

916 **Generalized Policy Iteration.** Sutton and Barto [55] consider most RL algorithms can be described in the
917 paradigm of Generalized Policy Iteration (GPI). In recent decade, RL algorithms usually resort to function
918 approximation (e.g., deep neural networks) to deal with large and continuous state space. An illustration of
919 GPI with function approximation is on the left of Figure 1. We use θ to denote the parameters of parameterized
920 value functions. Without loss of generality, we do not plot the parameters of policy since it is not necessary
921 for parameterized policy to exist, e.g., value-based RL algorithms [35]. For policy evaluation, value function
922 approximators are updated in finite times to approximate the true values (i.e., $V_\theta(s) \rightarrow v^\pi(s)$, $Q_\theta(s, a) \rightarrow$
923 $q^\pi(s, a)$), yet can never be perfect. For policy improvement, the policy are improved with respected to the
924 approximated value functions in an implicit (e.g., value-based RL) or explicit way (policy-based RL). In deep
925 RL, perfect policy evaluation and effective policy improvement are non-trivial to obtain with complex non-linear
926 function approximation from deep neural networks, thus most convergence and optimality results in conventional
927 RL usually no longer hold. From these two aspects, many works study how to improve the value function
928 approximation [61, 4, 27] and to propose more effective policy optimization or search algorithms [48, 50, 15].

929 E.2 A Unified View of Extensions of Conventional Value Function from the Vector Form of 930 Bellman Equation

931 Recall the vector form of Bellman equation (Equation 14), it indicates that value function is a function of policy
932 π and environmental models (i.e., \mathcal{P} and r). In conventional value functions and approximators, only state (and
933 action) is usually taken as input while other components in Equation 14 are modeled implicitly. Beyond state
934 (and action), consider explicit representation of some of components in Equation 14 during value estimation
935 can develop the ability of conventional value functions in different ways, to solve challenging problems, e.g.,
936 goal-conditioned RL [45, 1], Hierarchical RL [37, 64], opponent modeling and ad-hoc team [19, 14, 57], and
937 context-based Meta-RL [43, 29].

938 Most extensions of conventional VFA mentioned above are proposed for the purpose of value generalization
939 (among different space). Therefore, we suggest such extensions are derived from the same start point (i.e.,
940 Equation 14) and differ at the objective to represent and take as additional input explicitly of conventional value
941 functions. We provide a unified view of such extensions below:

- 942 • Goal-conditioned RL and context-based meta-RL usually focus on a series of tasks with similar goals
943 and environment models (i.e., \mathcal{P} and r). With goal representation as input, usually a subspace of
944 state space [45, 1], a value function approximation (VFA) can generalize values among goal space.
945 Similarly, with context representation [43, 11, 42], values generalize in meta tasks.

- 946 • Opponent modeling, ad-hoc team [19, 14, 57] seek to generalize among different opponents or
947 teammates in a Multiagent System, with learned representation of opponents. This can be viewed as
948 a special case of value generalization among environment models since from one agent view, other
949 opponents are part of the environment which also determines the dynamics and rewards. In multiagent
950 case, one can expand and decompose the corresponded joint policy in Equation 14 to see this.
- 951 • Hierarchical RL is also a special case of value generalization among environment models. In goal-
952 reaching fashioned Hierarchical HRL [37, 30, 38], high-level controllers (policy) issue goals for
953 low-level controls at an abstract temporal scale, while low-level controls take goals also as input and
954 aim to reach the goals. For low-level policies, a VFA with a given or learned goal representation space
955 can generalize values among different goals, similar to the goal-conditioned RL case as discussed
956 above. Another perspective is to view the separate learning process of hierarchical policies for different
957 levels as a multiagent learning system. Recently, a work [64] follows this view and extends high-level
958 policy with representation of low-level learning.

959 The common thing of above is that, they learn a representation of the environment (we call *external variables*).
960 In contrast, we study value generalization among agent’s own policies in this paper, which cares about *internal*
961 *variables*, i.e., the learning dynamics inside of the agent itself.

962 **Relation between PeVFA Value Approximation and Context-based Meta-RL.** For a given MDP, performing
963 a policy in the MDP actually induces a Markov Reward Process (MRP) [55]. One can view the policy and
964 actions are absorbed in the transition function of MRP. A value function defines the expected long-term returns
965 starting from a state. Therefore, different policies induces different MRPs and PeVFA value approximation
966 can be considered as a meta prediction task. In analogy to context-based Meta-RL where a task context is
967 learned to capture the underlying transition function and reward function of a MDP (i.e., task), one can view
968 policy representation as the context of corresponding MRP, since it is the underlying variable that determines the
969 transition function of MRPs.

970 E.3 A Review of Works on Policy Representation/Embedding Learning

971 Recent years, a few works involve representation or embedding learning for RL policy [18, 14, 3, 42, 64, 17].
972 We provide a brief review and summary for above works below.

973 The most common way to learn a policy representation is to extract from interaction trajectories through
974 policy recovery (i.e., behavioral cloning). For Multiagent Opponent Modeling [14], a policy representation
975 is learned from interaction episodes (i.e., state-action trajectories) through a *generative loss* and *discriminate*
976 *loss*. Generative loss is the same as the policy recovery auxiliary loss; discriminate loss is a triplet loss that
977 minimize the representation distance of the same policy and maximize those of different ones, which has the
978 similar idea of Contrastive Learning [41, 54]. Such opponent policy representations are used for prediction
979 of interaction outcomes for ad-hoc teams and are taken in policy network for some learning agent to facilitate
980 the learning when cooperating or competing with unknown opponents. More recently, in Hierarchical RL
981 [64], a representation is learned to model the low-level policy through *generative loss* mentioned above. The
982 low-level policy representation is taken in high-level policy to counter the non-stationarity issue of co-learning
983 of hierarchical policies. Later, Raileanu et al. [42] resort to almost the same method and the learned policy
984 representation is taken in their proposed PDVF. Along with a task context, the policy for a specific task can
985 be optimized in policy representation space, inducing a fast adaptation in new tasks. In summary, such a
986 representation learning paradigm can be considered as Surface Policy Representation (SPR) for policy data
987 encoding (trajectories as a special form of state-action pairs) plus policy recovery auxiliary loss (AUX) as we
988 introduced in Sec. 4.

989 A recent work [17] proposes Policy Evaluation Network (PVN) to approximate objective function $J(\pi)$. We
990 consider PVN as an predecessor of PDVF we mentioned above since offline policy optimization is also conducted
991 in learned representation space in a single task after similarly well training the PVN on many policies. The
992 authors propose *Network Fingerprint* to represent policy network. To circumvent the difficulty of representing the
993 parameters directly, policy network outputs (policy distribution) under a set of *probing states* are concatenated
994 and then taken as policy representation. Such probing states are randomly sampled for initialization and
995 also optimized with gradients through PVN and policies, like a search in joint state space. In principle, we
996 also consider this as a special instance of SPR, because network fingerprint follows the idea of reflecting the
997 information of how policy can behave under some states. Intuitively from a geometric view, this can be viewed
998 as using the concatenation of several representative (as denoted by the probing states) cross-sections in policy
999 surface (e.g., Figure 11) to represent a policy. For another view, one can imagine an equivalent case between
1000 SPR and network fingerprint, when state-action pairs of a deterministic policy are processed in SPR and a
1001 representation consisting of a number of actions under some key states or representative states is used in network
1002 fingerprint. Two potential issues may exist for network fingerprint. First, the dimensionality of representation
1003 is proportional to the number of probing states (i.e., $n|A|$), where a dilemma exists: more probing states are
1004 more representative while dimensionality can increase correspondingly. Second, it can be non-trivial and even

1005 unpractical to optimize probing states in the case with relatively state space of high dimension, which introduces
1006 additional computational complexity and optimization difficulty.

1007 In another concurrent work [10], a class of Parameter-based Value Functions (PVFs) are proposed. Instead of
1008 learning or designing a representation of policy, PVFs simply parse all the policy weights as inputs to the value
1009 function (i.e., Raw Policy Representation as also mentioned in our paper), even in the nonlinear case. Apparently,
1010 this can result in a unnecessarily large representation space which increase the difficulty of approximation and
1011 generalization. The issues of naively flattening the policy into a vector input are also pointed out in PVN [17].

1012 Others, several works in Policy Adaptation and Transfer [18, 3], Gaussian policy embedding representations are
1013 construct through Variational Inference.

1014 F Experimental Details and Complete Results

1015 F.1 Experimental Details

1016 **Environment.** We conduct our experiments on commonly adopted OpenAI Gym¹ continuous control tasks
1017 [6, 58]. We use the OpenAI Gym with version 0.9.1, the mujoco-py with version 0.5.4 and the MuJoCo products
1018 with version MJPRO131. Our codes are implemented with Python 3.6 and TensorFlow.

1019 **Resources and Equipment Used.** Our experiments are mainly conducted on a NVIDIA GeForce RTX 2080
1020 Ti with 11 GB memory. A single run of PPO-PeVFA usually takes 3-4 hours with about 6 trials are running
1021 simultaneously on the same GPU.

1022 **Implementation.** We use Proximal Policy Optimization (PPO) [50] with Generalized Advantage Estimator
1023 (GAE) [49] as our baseline algorithm. Recent works [8, 2] point out code-level optimizations influence the
1024 performance of PPO a lot. For a fair comparison and clear evaluation, we perform no code-level optimization
1025 in our experiments, e.g., state standardization, reward scaling, gradient clipping, parameter sharing and etc.
1026 Our proposed algorithm PPO-PeVFA is implemented based on PPO, which only differs at the replacement
1027 for conventional value function network with PeVFA network. Policy network is a 2-layer MLP with 64
1028 units per layer and ReLU activation, outputting a Gaussian policy, i.e., a tanh-activated mean along with a
1029 state-independent vector-parameterized log standard deviation. For PPO, the conventional value network $V_\phi(s)$
1030 (VFA) is a 2-layer 128-unit ReLU-activated MLP with state as input and value as output. For PPO-PeVFA, the
1031 PeVFA network $\mathbb{V}_\theta(s, \chi_\pi)$ takes as input state and policy representation χ_π which has the dimensionality of 64,
1032 with the structure illustrated in Figure 8(a). We do not use parameter sharing between policy and value function
1033 approximators for more clear evaluation.

1034 **Training and Evaluation.** We use Monte Carlo returns for value approximation. In contrast to conventional
1035 VFA V_ϕ which approximates the value of current policy (e.g., Algorithm 2), PeVFA $\mathbb{V}_\theta(s, \chi_\pi)$ is additionally
1036 trained to approximate the values of all historical policies ($\{\pi_i\}_{i=0}^t$) along the policy improvement path (e.g.,
1037 Algorithm 3). The policy network parameterized by ω is then updated with following loss function:

$$\mathcal{L}^{\text{PPO}}(\omega) = -\mathbb{E}_{\pi_{\omega^-}} \left[\min \left(\rho_t \hat{A}(s_t, a_t), \text{clip}(\rho_t, 1 - \epsilon, 1 + \epsilon) \hat{A}(s_t, a_t) \right) \right], \quad (15)$$

1038 where $\hat{A}(s_t, a_t)$ is advantage estimation of old policy π_{ω^-} , which is calculated by GAE based on conventional
1039 VFA V_ϕ or PeVFA $\mathbb{V}_\theta(s, \chi_\pi)$ respectively, and $\rho_t = \frac{\pi_\omega(a_t, s_t)}{\pi_{\omega^-}(a_t, s_t)}$ is the importance sampling ratio. Note that
1040 both PPO and PPO-PeVFA update the policy according to Equation 15 and only differ at advantage estimation
1041 based on conventional VFA V_ϕ or PeVFA $\mathbb{V}_\theta(s, \chi_\pi)$. This ensures that the performance difference comes only
1042 from different approximation of policy values. Common learning parameters for PPO and PPO-PeVFA are
1043 shown in Table 2. For each iteration, we update value function approximators first and then the policy with
1044 updated values. Such a training scheme is used for both PPO and PPO-PeVFA. For evaluation, we evaluate
1045 the learning algorithm every 20k time steps, averaging the returns of 10 episodes. Fewer evaluation points are
1046 selected and smoothed over neighbors and then plotted in our learning curves below.

1047 **Details for PPO-PeVFA.** For PeVFA, the training process also involves value approximation of historical
1048 policies and learning of policy representation. Training details are shown in Table 3. PeVFA $\mathbb{V}_\theta(s, \chi_\pi)$ is
1049 trained every 10 steps with a batch of 64 samples from an experience buffer with size of 200k steps. Policy
1050 representation model is trained at intervals of 10 or 20 steps depending on OPR or SPR adopted. Due to 1k - 2k
1051 policies are collected in total in each trial, a relatively small batch size of policy is used. For OPR, Random
1052 Mask (Figure 13) is performed on all weights and biases of policy network except for the output layer (i.e., mean
1053 and log-std). For SPR, two buffers of state-action pairs are maintained for each policy: a small one is sampled
1054 for calculating SPR and the relatively larger one is sampled for auxiliary training (policy recovery).

¹<http://gym.openai.com/>

Table 2: Common hyperparameter choices of PPO and PPO-PeVFA.

| Hyperparameters for PPO & PPO-PeVFA | |
|-----------------------------------------|-------------------------------|
| Policy Learning Rate | 10^{-4} |
| Value Learning Rate | 10^{-3} |
| Clipping Range Parameter (ϵ) | 0.2 |
| GAE Parameter (λ) | 0.95 |
| Discount Factor (γ) | 0.99 |
| On-policy Samples Per Iteration | 5 episodes or 2000 time steps |
| Batch Size | 128 |
| Actor Epoch | 10 |
| Critic Epoch | 10 |
| Optimizer | Adam |

Table 3: Training details for PPO-PeVFA, including value approximation of historical policies and policy representation learning. CL is abbreviation for Contrastive Learning and AUX is for auxiliary loss of policy recovery. In our experiments, grid search is performed for the best hyperparameter configuration regarding terms with multiple alternatives (i.e., { }).

| Value Approximation for Historical Policies | |
|---------------------------------------------|-------------------------------------------------------|
| Value Learning Rate | 10^{-3} |
| Training Frequency | Every 10 time steps |
| Batch Size | 64 |
| Experience Buffer Size | 200k (steps) |
| Policy Representation Learning | |
| Training Frequency | Every {10, 20} time steps |
| Policy Num Per Batch | {16, 32} |
| SPR s, a Pair Num | {200, 500} |
| CL Learning Rate | { 10^{-3} , $5 \cdot 10^{-4}$ } |
| CL Momentum | { $5 \cdot 10^{-2}$, 10^{-2} , $5 \cdot 10^{-3}$ } |
| CL Mask Ratio for OPR | {0.1, 0.2} |
| CL Sample Ratio for SPR | 0.8 |
| AUX Learning Rate | 10^{-3} |
| AUX Batch Size | {128, 256} |

1055 **F.2 Complete Learning Curves for Evaluation Results in Table 1**

1056 Corresponding to Table 1, an overall view of learning curves of all variants of PPO-PeVFA as well as baseline
 1057 algorithms are shown in Figure 17. One can refer to Figure 14 for a clearer view of the effects of PeVFA (with
 1058 E2E policy representation), and Figure 15, 16 for the effects of self-supervised policy representations, i.e., CL
 1059 and AUX.

1060 **F.3 Visualization of Learned Policy Representation**

1061 To show how the learned representation is like in a low-dimensional space, we visualize the learned representation
 1062 of policies encountered during the learning process.

1063 **Visualization Design.** We record all policies on the policy improvement path during the learning process of a
 1064 PPO-PeVFA agent. For each trial in our experiments in MuJoCo continuous control tasks, about 1k - 2k policies
 1065 are collected. We run 5 trials and 5k - 12k policies are collected in total for each task. We also store the policy
 1066 representation model at intervals for each trial, and we use last three checkpoints to compute the representation
 1067 of each policy collected. For each policy collected during 5 trials, its representation for visualization is obtained
 1068 by averaging the results of 3 checkpoints of each trial and then concatenating the results from 5 trials. Finally,
 1069 we plot 2D embedding of policy representations prepared above through t-SNE [34] and Principal Component
 1070 Analysis (PCA) in `sklearn`².

²<https://scikit-learn.org/stable/index.html>

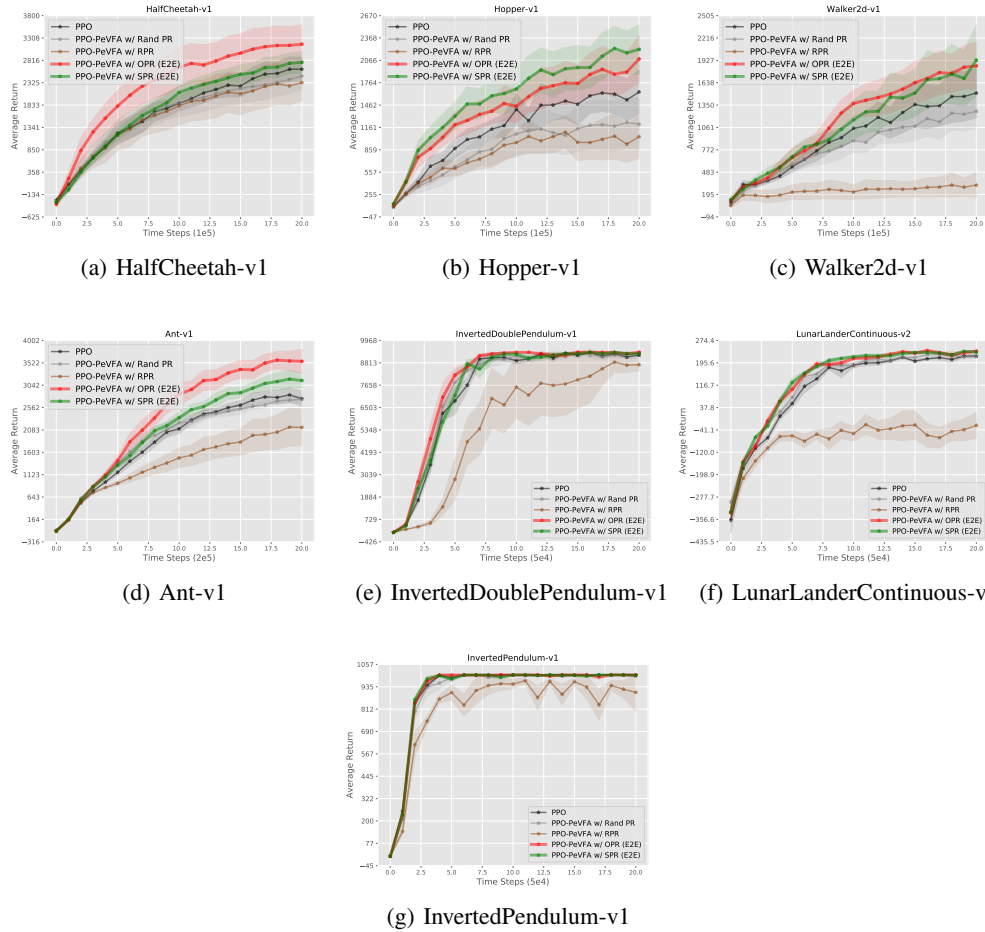


Figure 14: Evaluations of PPO-PeVFA with end-to-end (E2E) trained OPR and SPR in MuJoCo continuous control tasks. The results demonstrate the effectiveness of PeVFA and two kinds of policy representation, answering the Question 1. The results are average returns and the shaded region denotes half a standard deviation over 10 trials.

1071 **Results and Analysis.** Visualizations of OPR and SPR learned in an end-to-end fashion in HalfCheetah-v1 and
 1072 Ant-v1 are in Figure 18 and 19. We use different types of markers to distinguish policies from different trials to
 1073 see how policy evolves in representation space from different random initialization. Moreover, we provide two
 1074 views: performance view and process view, to see how policies are aligned in representation space regarding
 1075 performance and ‘age’ of policies respectively.

1076 Visualization of OPR trained in end-to-end fashion is shown in Figure 18. From the performance view, it is
 1077 obvious that policies of poor and good performances are aligned from left to right in t-SNE representation space
 1078 and are aligned at two distinct directions in PCA representation space. An evolution of policies from different
 1079 trials can be observed in subplot (b) and (d). Thus, policies from different trials are locally continuous; while
 1080 policies are globally consistent in representation space with respect to policy performance. Moreover, we can
 1081 observe multimodality for policies with comparable performance. This means that the obtained representation
 1082 not only reflects optimality information but also maintains the behavioral characteristic of policy.

1083 Parallel to OPR, end-to-end trained SPR is visualized in Figure 19. A more obvious multimodality can be
 1084 observed in both t-SNE and PCA space: policies from different trials start from the same region and then
 1085 diverge during the following learning process. Different from OPR, SPR shows more distinction among different
 1086 trials since SPR is a more direct reflection of policy behavior (*dynamics* property as mentioned in Sec. D.5).
 1087 Another thing is, policies from different trials forms wide ‘strands’ especially in t-SNE representation space.
 1088 We conjecture that it is because SPR is a more stochastic way to obtain representation as random selected
 1089 state-action pairs are used.

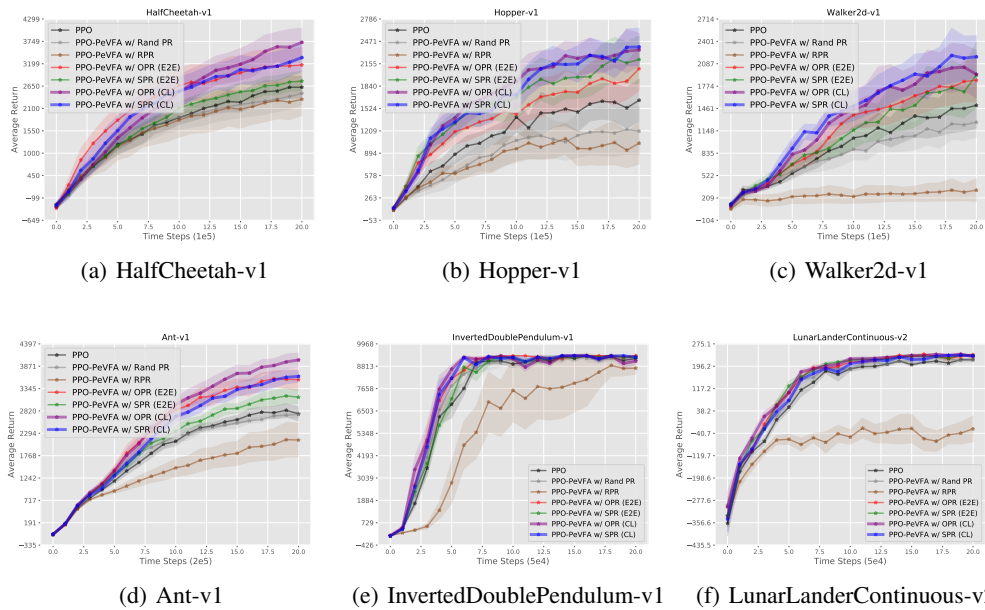


Figure 15: Evaluations of PPO-PeVFA with OPR and SPR trained through contrastive learning (CL) in MuJoCo continuous control tasks. The results are average returns and the shaded region denotes half a standard deviation over 10 trials.

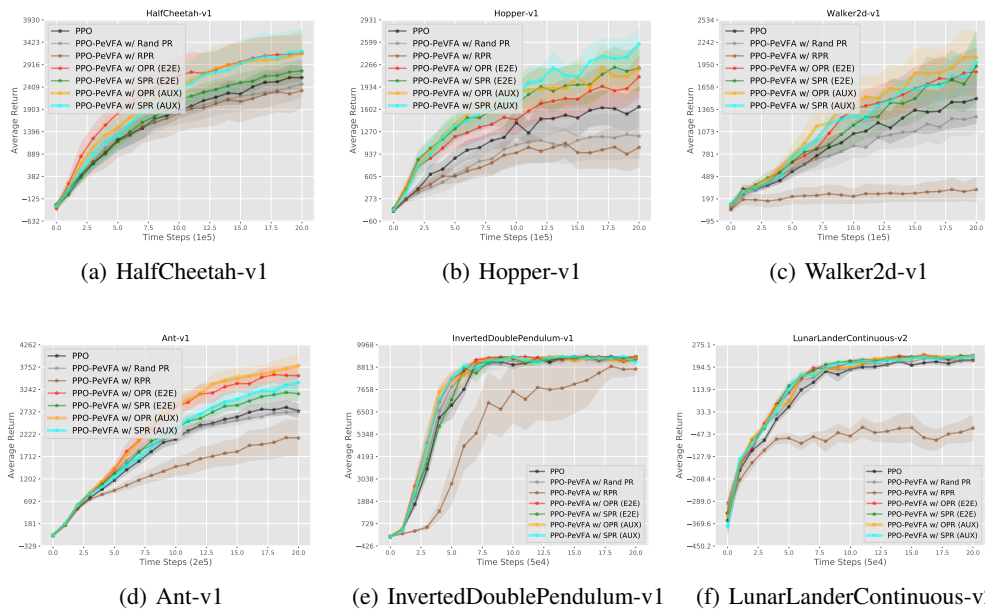
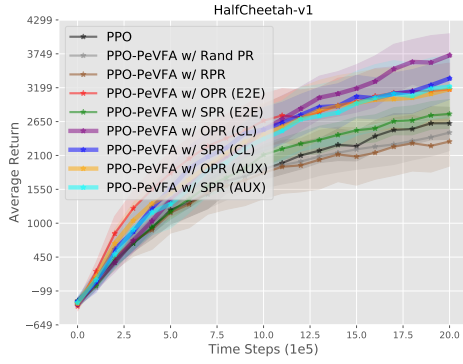
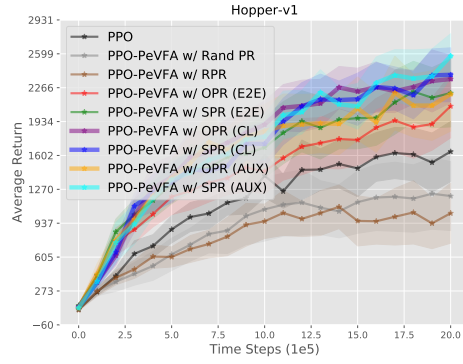


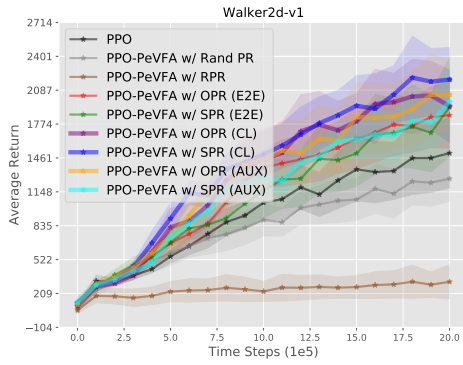
Figure 16: Evaluations of PPO-PeVFA with OPR and SPR trained through auxiliary loss of policy recovery (AUX) in MuJoCo continuous control tasks. The results are average returns and the shaded region denotes half a standard deviation over 10 trials.




(a) HalfCheetah-v1



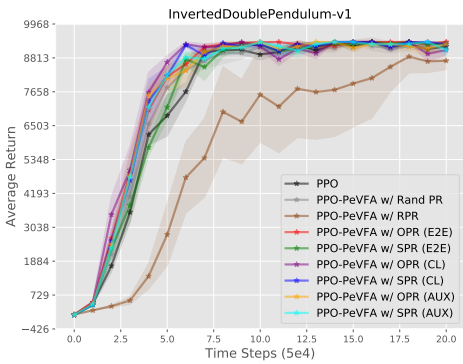
(b) Hopper-v1



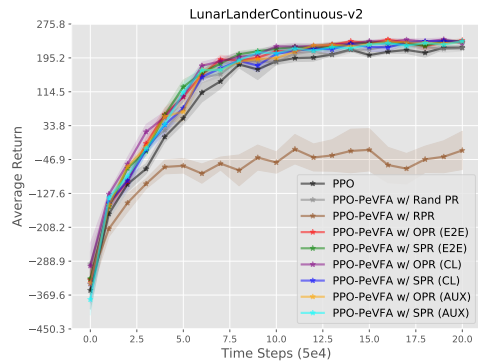
(c) Walker2d-v1



(d) Ant-v1



(e) InvertedDoublePendulum-v1



(f) LunarLanderContinuous-v2

Figure 17: An overall view of performance evaluations of different algorithms in MuJoCo continuous control tasks. The results are average returns and the shaded region denotes half a standard deviation over 10 trials.

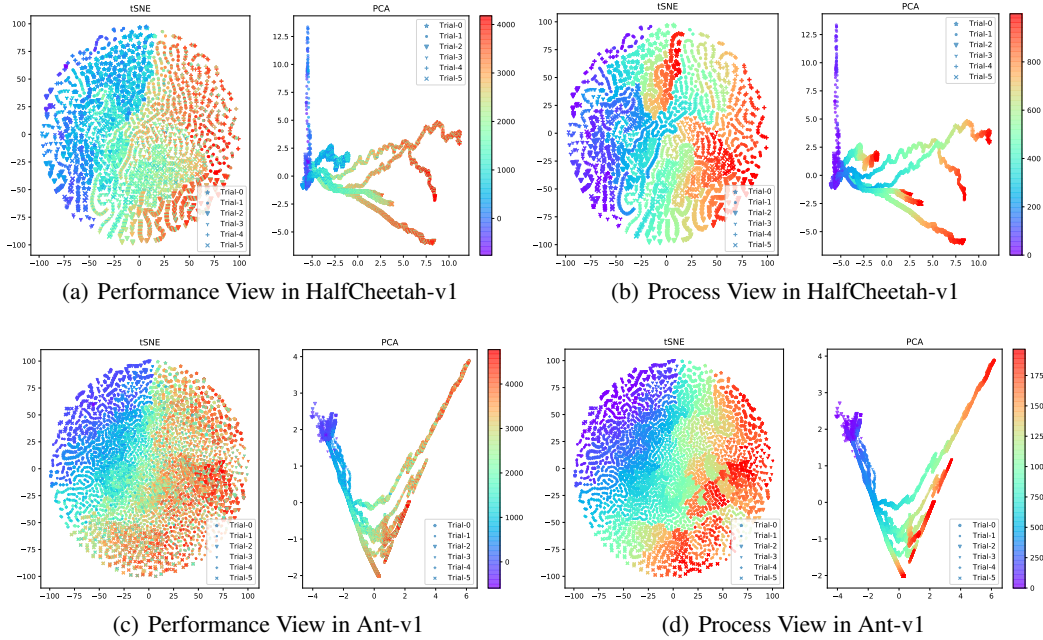


Figure 18: Visualizations of end-to-end (E2E) learned Origin Policy Representation (OPR) for policies collected during 5 trials (denoted by different kinds of markers). In total, about 6k policies are plotted for HalfCheetah-v1 (a-b) and 12k for Ant-v1 (c-d). In each subplot, t-SNE and PCA 2D embeddings are at left and right respectively. In performance view, each policy (i.e., marker) is colored by its performance evaluation (averaged return). In process view, each policy is colored by its corresponding iteration ID during GPI process.

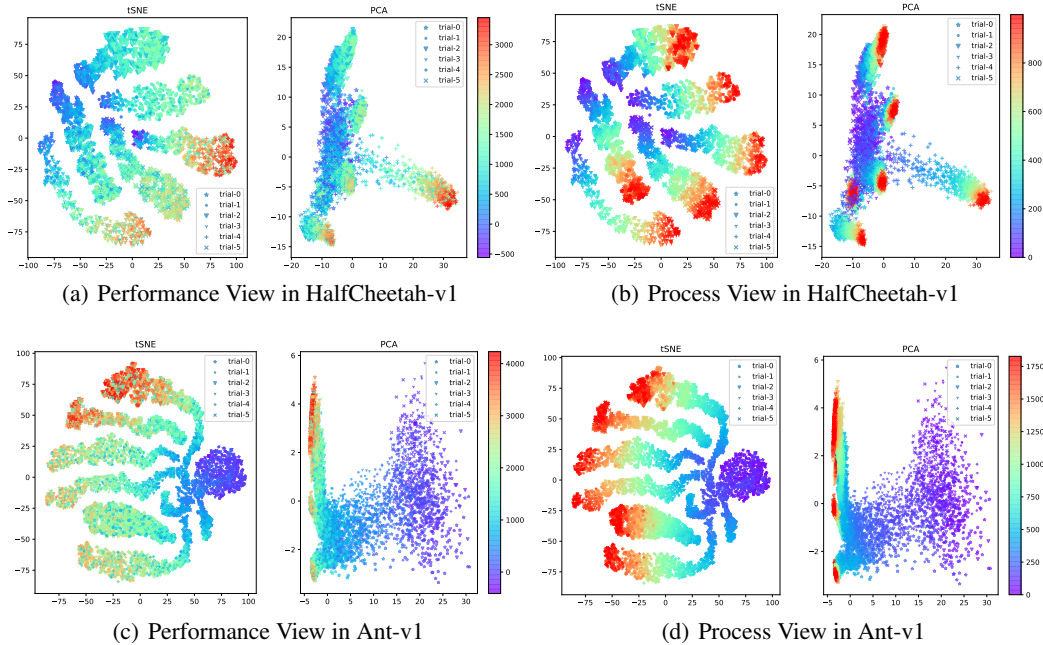


Figure 19: Visualizations of end-to-end (E2E) learned Surface Policy Representation (SPR) for policies collected during 5 trials (denoted by different kinds of markers). In performance view, each policy (i.e., marker) is colored by its performance evaluation (averaged return). In process view, each policy is colored by its corresponding iteration ID during GPI process.