
Iteratively Learn Diverse Strategies with State Distance Information

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In complex reinforcement learning (RL) problems, policies with similar rewards
2 may have substantially different behaviors. It remains a fundamental challenge
3 to optimize rewards while also discovering as many *diverse* strategies as possible,
4 which can be crucial in many practical applications. Our study examines two design
5 choices for tackling this challenge, i.e., *diversity measure* and *computation frame-*
6 *work*. First, we find that with existing diversity measures, visually indistinguishable
7 policies can still yield high diversity scores. To accurately capture the behavioral
8 difference, we propose to incorporate the state-space distance information into
9 the diversity measure. In addition, we examine two common computation frame-
10 works for this problem, i.e., population-based training (PBT) and iterative learning
11 (ITR). We show that although PBT is the precise problem formulation, ITR can
12 achieve comparable diversity scores with higher computation efficiency, leading to
13 improved solution quality in practice. Based on our analysis, we further combine
14 ITR with two tractable realizations of the state-distance-based diversity measures
15 and develop a novel diversity-driven RL algorithm, *State-based Intrinsic-reward*
16 *Policy Optimization* (SIPO), with provable convergence properties. We empirically
17 examine SIPO across three domains from robot locomotion to multi-agent games.
18 In all of our testing environments, SIPO consistently produces strategically diverse
19 and human-interpretable policies that cannot be discovered by existing baselines.

20 1 Introduction

21 A consensus in deep learning (DL) is that different local optima have similar mappings in the
22 functional space, leading to similar losses to the global optimum [57, 52, 36]. Hence, via stochastic
23 gradient descent (SGD), most DL works only focus on the final performance without considering
24 *which* local optimum SGD discovers. However, in complex reinforcement learning (RL) problems, the
25 policies associated with different local optima can exhibit significantly different behaviors [9, 31, 59].
26 Thus, it is a fundamental problem for an RL algorithm to not only optimize rewards but also discover
27 as many diverse strategies as possible. A pool of diversified policies can be further leveraged
28 towards a wide range of applications, including the discovery of emergent behaviors [30, 2, 56],
29 generating diverse dialogues [26], designing robust robots [11, 22, 17], and enhancing human-AI
30 collaboration [34, 8, 10].

31 Obtaining diverse RL strategies requires a quantitative method for measuring the difference (i.e.,
32 *diversity*) between two policies. However, how to define such a measure remains an open challenge.
33 Previous studies have proposed various diversity measures, such as comparing the difference between
34 the action distributions generated by policies [55, 34, 69], computing probabilistic distances between
35 the state occupancy of different policies [39], or measuring the mutual information between states and
36 policy identities [14]. However, it remains unclear which measure could produce the best empirical
37 performance. Besides, the potential pitfalls of these measures are rarely discussed.

38 In addition to diversity measures, there are two common computation frameworks for discovering
39 diverse policies, including population-based training (PBT) and iterative learning (ITR). PBT directly
40 solves a constrained optimization problem by learning a collection of policies simultaneously, subject
41 to policy diversity constraints [48, 34, 8]. Although PBT is perhaps the most popular framework in
42 the existing literature, it can be computationally challenging [44] since the number of constraints
43 grows quadratically with the number of policies. The alternative framework is ITR, which iteratively
44 learns a single policy that is sufficiently different from previous policies [39, 69]. ITR is a greedy
45 relaxation of the PBT framework and it largely simplifies the optimization problem in each iteration.
46 However, the performance of the ITR framework has not been theoretically analyzed yet, and it is
47 often believed that ITR can be less efficient due to its sequential nature.

48 We provide a comprehensive study of the two aforementioned design choices. First, we examine
49 the limitations of existing diversity measures in a few representative scenarios, where two policies
50 outputting very different action distributions can still lead to similar state transitions. In these
51 scenarios, state-occupancy-based measures are not sufficient to truly reflect the underlying behavior
52 differences of the policies either. By contrast, we observe that diversity measures based on *state*
53 *distances* can accurately capture the visual behavior differences of different policies. Therefore, we
54 suggest that an effective diversity measure should explicitly incorporate state distance information
55 for the best practical use. Furthermore, for the choice of computation framework, we conduct an
56 in-depth analysis of PBT and ITR. We provide theoretical evidence that ITR, which has a simplified
57 optimization process with fewer constraints, can discover solutions with the same reward as PBT
58 while achieving *at least half* of the diversity score. This finding implies that although ITR is a greedy
59 relaxation of PBT, their optimal solutions can indeed have comparable qualities. Furthermore, note
60 that policy optimization is much simplified in ITR, which suggests that ITR can result in much better
61 empirical performances and should be preferred in practice.

62 Following our insights, we combine ITR and a state-distance-based diversity measure to develop a
63 generic and effective algorithm, *State-based Intrinsic-reward Policy Optimization (SIPO)*, for discover-
64 ing diverse RL strategies. In each iteration, we further solve this constrained optimization problem
65 via Lagrangian method and two-timescale gradient descent ascent (GDA) [27]. We theoretically
66 prove that our algorithm is guaranteed to converge to a neighbor of ϵ -stationary point. Regarding the
67 diversity measure, we provide two practical realizations, including a straightforward version based
68 on the RBF kernel and a more general learning-based variant using Wasserstein distance.

69 We evaluate SIPO in three domains ranging from single-agent continuous control to multi-agent
70 games: Humanoid locomotion [38], StarCraft Multi-Agent Challenge [53], and Google Research
71 Football (GRF) [23]. Our findings demonstrate that SIPO surpasses baselines in terms of population
72 diversity score across all three domains. Remarkably, our algorithm can successfully discover 6
73 distinct human-interpretable strategies in the GRF 3-vs-1 scenario and 4 strategies in two 11-player
74 GRF scenarios, namely counter-attack and corner, without any domain-specific priors, which are
75 beyond the capabilities of existing algorithms.

76 2 Related Work

77 **Diversity in RL.** It has been shown that policies trained under the same reward function can exhibit
78 significantly different behaviors [9, 31]. Merely discovering a single high-performing solution may
79 not suffice in various applications [11, 59, 22]. As such, the discovery of a diverse range of policies
80 is a fundamental research problem, garnering attention over many years [40, 12, 24]. Early works
81 are primarily based on multi-objective optimization [41, 51, 35, 43, 50], which assumes a set of
82 reward functions is given in advance. In RL, this is also related to reward shaping [42, 1, 13, 56]. We
83 consider learning diverse policies without any domain knowledge.

84 **Population-based training (PBT)** is the most popular framework for diverse solutions by jointly
85 learning separate policies. Representative works include evolutionary computation [60, 33, 48],
86 league training [59, 20], computing Hessian matrix [47] or constrained optimization with a population
87 diversity measure [34, 68, 25, 32, 8]. An improvement is to learn a latent variable policy instead of
88 separate ones. Prior works have incorporated different domain knowledge to design the latent code,
89 such as action clustering [61], agent identities [25] or prosocial level [49, 2]. The latent variable can
90 be also learned in an unsupervised fashion, such as in DIYAN [14] and its variants [22, 45]. Zahavy
91 et al. [66] learns latent-conditioned diverse policies with hard constraints on rewards to ensure the

92 derived policies are (nearly) optimal. In contrast, we prioritize diversity and fully accept sub-optimal
 93 strategies, leading to a hard constraint on diversity measures.

94 **Iterative learning (ITR)** simplifies PBT by only optimizing a single policy in each iteration and
 95 forcing it to behave differently w.r.t. previously learned ones [39, 55, 69]. While some ITR works
 96 require an expensive clustering process before each iteration [67] or domain-specific features [65], we
 97 consider domain-agnostic ITR in an end-to-end fashion. Besides, Pacchiano et al. [46] learns a kernel-
 98 based score function to iteratively guide policy optimization. The score function is conceptually
 99 similar to SIPO-WD but is applied to a parallel setting with more restricted expressiveness power.

100 **Diversity Measure.** Most previous works considered diversity measures on action distribution
 101 and state occupancy. For example, measures such as Jensen-Shannon divergence [34] and cross-
 102 entropy [69] are defined over policy distributions to encourage different policies to take different
 103 actions on the same state, implicitly promoting the generation of diverse trajectories. Other measures
 104 such as maximum mean discrepancy [39] maximize the probability distance between the state
 105 distributions induced by two policies. However, these approaches can fail to capture meaningful
 106 behavior differences between two policies in certain scenarios, as we will discuss in Section 4.1. There
 107 also exist specialized measures, such as cross-play rewards [8], which is designed for cooperative
 108 multi-agent games. It is worth noting that diversity measures are closely related to exploration
 109 criteria [3, 18, 5] and skill discovery [7, 28, 21], where a diversity surrogate objective is often
 110 introduced to encourage broad state coverage. However, this paper aims to explicitly discover
 111 mutually distinct policies. Our diversity measure depends on a function that computes the distance
 112 between states visited by two policies.

113 3 Preliminary

114 **Notation:** We consider POMDP [54] defined by $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{O}, r, P, O, \nu, H \rangle$. \mathcal{S} is the state
 115 space. \mathcal{A} and \mathcal{O} are the action and observation space. $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function.
 116 $O : \mathcal{S} \rightarrow \mathcal{O}$ is the observation function. H is the horizon. P is the transition function. At timestep
 117 h , the agent receives an observation $o_h = O(s_h)$ and outputs an action $a_h \in \mathcal{A}$ w.r.t. its policy
 118 $\pi : \mathcal{O} \rightarrow \Delta(\mathcal{A})$. The RL objective $J(\pi)$ is defined by $J(\pi) = \mathbb{E}_{(s_h, a_h) \sim (P, \pi)} \left[\sum_{h=1}^H r(s_h, a_h) \right]$.
 119 The above formulation can be naturally extended to cooperative multi-agent settings, where π and
 120 R correspond to the joint policy and the shared reward. Moreover, in this paper, we **assume access**
 121 **to object-centric information and features** rather than pure visual observations to simplify our
 122 discussion. We remark that although we restrict the scope of this paper to states, our method can be
 123 further extended to high-dimensional inputs (e.g. images, see App. B.3) via representation learning.
 124 Finally, to discover diverse strategies, we aim to learn a set of M policies $\{\pi_i\}_{i=1}^M$ such that all of
 125 these policies are locally optimal under $J(\cdot)$ but mutually distinct subject to some diversity measure
 126 $D(\cdot, \cdot) : \Delta \times \Delta \rightarrow \mathbb{R}$, which captures the difference between two policies.

127 **Existing Diversity Measures:** We say a diversity measure D is defined over action distribution if it
 128 can be written as

$$D(\pi_i, \pi_j) = \mathbb{E}_{s \sim q(s)} \left[\tilde{D}_{\mathcal{A}}(\pi_i(\cdot | s) \| \pi_j(\cdot | s)) \right], \quad (1)$$

130 where q is an occupancy measure over states, $\tilde{D}_{\mathcal{A}} : \Delta \times \Delta \rightarrow \mathbb{R}$ measures the difference between
 131 action distributions. $\tilde{D}_{\mathcal{A}}$ can be any probability distance as defined in prior works [55, 34, 69, 48].

132 Denote the state occupancy of π as q_{π} . We say a diversity measure is defined over state occupancy if
 133 it can be written as

$$D(\pi_i, \pi_j) = \tilde{D}_{\mathcal{S}}(q_{\pi_i} \| q_{\pi_j}), \quad (2)$$

134 which can be realized as an integral probability metric [39]. We remark that q_{π} is usually intractable.

136 In addition to diversity measures, we present two popular computation frameworks for this purpose.

137 **Population-Based Training (PBT):** PBT is a straightforward formulation by jointly learning M
 138 policies $\{\pi_i\}_{i=1}^M$ subject to pairwise diversity constraints, i.e.,

$$\max_{\{\pi_i\}} \sum_{i=1}^M J(\pi_i) \quad \text{s.t.} \quad D(\pi_j, \pi_k) \geq \delta, \forall j, k \in [M], j \neq k, \quad (3)$$

139 where δ is a threshold. In our paper, we consistently refer to the aforementioned computation
 141 framework as "PBT", rather than adjusting hyperparameters [19]. Despite a precise formulation, PBT
 142 poses severe optimization challenges due to mutual constraints.

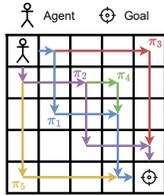


Table 1: Diversity measures of the grid-world example. Computation details can be found in App. B.

| | human | action-based | | | state-distance-based | | |
|-------------------|-------|--------------|------------------|-----------------------|----------------------|-------------|-------------|
| | | KL | JSD ₁ | JSD ₀ /EMD | L_2 norm | L_2 norm | EMD |
| $D(\pi_1, \pi_2)$ | small | $+\infty$ | $\log 2$ | $1/2$ | $\sqrt{7}$ | $2\sqrt{2}$ | 5.7 |
| $D(\pi_1, \pi_3)$ | large | $+\infty$ | $\log 2$ | $1/8$ | 1 | $2\sqrt{6}$ | 11.3 |

Figure 1: (left) A grid-world environment with 5 different optimal policies. Intuitively, $D(\pi_1, \pi_2) < D(\pi_1, \pi_3)$ and $D(\pi_3, \pi_4) < D(\pi_3, \pi_5)$. However, action-based measures can give $D_A(\pi_1, \pi_2) \geq D_A(\pi_1, \pi_3)$ and state-occupancy-based measures can give $D(\pi_3, \pi_4) = D(\pi_3, \pi_5)$.

143 **Iterative Learning (ITR):** ITR is a greedy approximation of PBT by iteratively learning novel
 144 policies. In the i -th ($1 \leq i \leq M$) iteration, ITR solves

$$\pi_i^* = \arg \max_{\pi_i} J(\pi_i) \text{ s.t. } D(\pi_i, \pi_j^*) \geq \delta, \forall 1 \leq j < i. \quad (4)$$

145 π_j^* is recursively defined by the above equation. Compared with PBT, ITR trades off wall-clock time
 146 for less required computation resources (e.g., GPU memory) and performs open-ended training (i.e.,
 147 the population size M does not need to be fixed at the beginning of training).

148 4 Analysis of Existing Diversity-Discovery Approaches

149 In this section, we conduct both quantitative and theoretical analyses of existing approaches to
 150 motivate our method. We first discuss diversity measures in Sec. 4.1 and then compare computation
 151 frameworks, namely PBT and ITR, in Sec. 4.2.

152 4.1 A Common Missing Piece in Diversity Measure: State Distance

153 The perception of diversity among humans primarily relies on the level of dissimilarity within the
 154 state space, which is measured by a distance function. However, the diversity measures outlined
 155 in Eq. (1) and Eq. (2) completely fail to account for such crucial information. In this section, we
 156 provide a detailed analysis to instantiate this observation with concrete examples and propose a novel
 157 diversity measure defined over state distances.

158 First, we present a synthetic example to demonstrate the limitations of current diversity measures.
 159 Our example consists of a grid-world environment with a single agent and grid size N_G . The agent
 160 starts at the top left of the grid-world and must navigate to the bottom right corner, as shown in
 161 Fig. 1. While N_G can be large in general, we illustrate with $N_G = 5$ for simplicity. We draw five
 162 distinct policies, denoted as π_1 through π_5 , which differ in their approach to navigating the grid-world.
 163 Consider π_1, π_2 , and π_3 first. Although humans may intuitively perceive that policies π_1 and π_2 ,
 164 which move along the diagonal, are more similar to each other than to π_3 , which moves along the
 165 boundary, diversity measures based on actions can fail to reflect this intuition, as shown in Table 1.
 166 Then, let's switch to policies π_3, π_4 , and π_5 . We find that state-occupancy-based diversity measures
 167 are unable to differentiate between π_4 and π_5 in contrast to π_3 . This is because the states visited by
 168 π_3 are entirely disjoint from those visited by both π_4 and π_5 . However, humans would judge π_5 to be
 169 more distinct from π_3 than π_4 because both π_3 and π_4 tend to visit the upper boundary.

170 Next, we consider a more realistic and complicated multi-agent football scenario in Fig. 2, where an
 171 idle player in the backyard takes an arbitrary action without involving in the attack at all. Although
 172 the idle player stays still with no effect on the team strategy, action-based measures can produce high
 173 diversity scores when the idle player takes different duplicated actions.

174 To summarize, existing measures suffer from a significant limitation — they only compare the
 175 behavior trajectories *implicitly* through the lens of action or state distribution without *explicitly*
 176 *measuring state distance*. Specifically, action-based measures fail to capture the behavioral differences
 177 that may arise when similar states are reached via different actions. Similarly, state occupancy
 178 measures do not quantify *the degree of dissimilarity* between states. To address this limitation, we
 179 propose a new diversity measure that explicitly takes into account the distance function in state space:

$$D(\pi_i, \pi_j) = \mathbb{E}_{(s, s') \sim \gamma} [g(d(s, s'))], \quad (5)$$



Figure 2: Duplicate actions in multi-agent football. For players who are not involved in the attack, actions like “pass”, “shoot”, and “slide” result in the same consequence. Diversity measures should not focus on these actions.

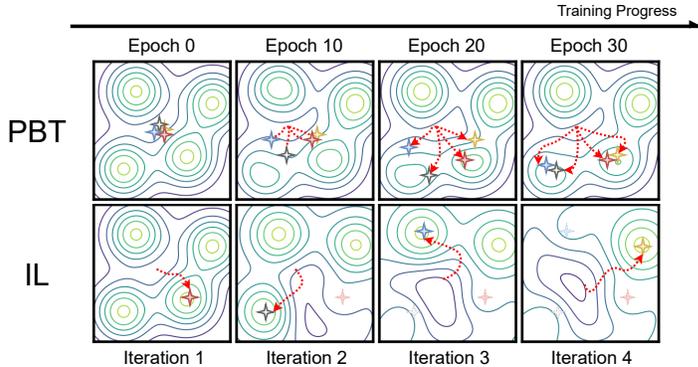


Figure 3: Illustration of the learning process of PBT and ITR in a 2-D navigation environment with 4 modes. PBT will not uniformly converge to different landmarks as computation can be either too costly or unstable. By contrast, ITR repeatedly excludes a particular landmark, such that policy in the next iteration can continuously explore until a novel landmark is discovered.

180 d is a distance metric over $\mathcal{S} \times \mathcal{S}$. $g : \mathbb{R}^+ \rightarrow \mathbb{R}$ is a monotonic cost function. $\gamma \in \Gamma(q_{\pi_i}, q_{\pi_j})$
 181 is a distribution over state pairs. $\Gamma(q_{\pi_i}, q_{\pi_j})$ denotes the collection of all distributions on $\mathcal{S} \times \mathcal{S}$
 182 with marginals q_{π_i} and q_{π_j} on the first and second factors respectively. We also note that states are
 183 consequences of performed actions. Hence, a state-distance-based measure also implicitly reflects the
 184 (meaningful) differences in actions between two policies. We compute two simple measures based on
 185 state distance, i.e., the L_2 norm and the Earth Moving Distance (EMD), for the grid-world example
 186 and present results in Table 1. These measures are consistent with human intuition.

187 4.2 Computation Framework: Population-Based or Iterative Learning?

188 We first consider a simplest motivating example to intuitively illustrate the optimization chal-
 189 lenges. Let’s assume that π_i is a scalar, $J(\pi_i)$
 190 is linear in π_i , and $D(\pi_i, \pi_j) = |\pi_i - \pi_j|$. In
 191 our definition, where M denotes the number
 192 of diverse policies, PBT involves $\Theta(M^2)$
 193 constraints in a single linear programming
 194 problem while ITR involves $\mathcal{O}(M)$ constraints in each of
 195 M iterations. Given that the complexity of linear
 196 programming is a high-degree polynomial
 197 (higher than 2) of the number of constraints,
 198 solving PBT is harder (and probably slower) than solving ITR in a total of M iterations, *despite PBT*
 199 *being parallelized*. This challenge can be more severe in RL due to complex solution space and large
 200 training variance.
 201

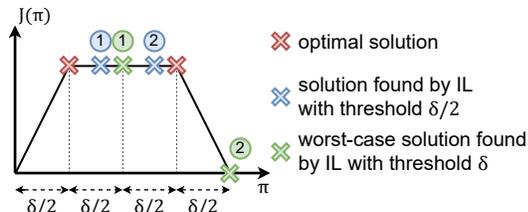


Figure 4: 1-D worst case of IL. With threshold δ , IL finds solutions with inferior rewards. However, IL can find optimal solutions if the threshold is halved.

202 Although ITR can be optimized efficiently, it remains unclear whether ITR, as a greedy approximation
 203 of PBT, can obtain solutions of comparable rewards. Fig. 4 shows the worst case in the 1-D setting
 204 when the ITR solutions (green) can indeed have lower rewards than the PBT solution (red) subject to
 205 the same diversity constraint. However, we will show in the next theorem that ITR is guaranteed to
 206 have no worse rewards than PBT by trading off half of the diversity.

207 **Theorem 4.1.** Assume D is a distance metric. Denote the optimal value of Eq.(3) as T_1 . Let
 208 $T_2 = \sum_{i=1}^M J(\tilde{\pi}_i)$ where

$$\tilde{\pi}_i = \arg \max_{\pi_i} J(\pi_i) \text{ s.t. } D(\pi_i, \tilde{\pi}_j) \geq \delta/2, \forall 1 \leq j < i \quad (6)$$

209 for $i = 1, \dots, M$, then $T_2 \geq T_1$.

210 Please see App. E.1 for the proof. The above theorem provides a quality guarantee for ITR. The proof
 211 can be intuitively explained by the 1-D example in Fig. 4, where green points represent the worst
 212 case with threshold δ and blue points represent the solutions with threshold $\delta/2$. Thm. 4.1 shows that,

213 for any policy pool derived by PBT, we can always use ITR to obtain another policy pool, which has
 214 *the same rewards and comparable diversity scores.*

215 **Empirical Results:** We empirically compare PBT and ITR in a
 216 2-D navigation environment with 1 agent and N_L landmarks in
 217 Fig. 3. The reward is 1 if the agent successfully navigates to a
 218 landmark and 0 otherwise. We train N_L policies using both PBT
 219 and ITR to discover strategies toward each of these landmarks.
 220 More details can be found in App. D. Table 2 shows the number
 221 of discovered landmarks by PBT and ITR. ITR performs consistently
 222 better than PBT even in this simple example. We intuitively
 223 illustrate the learning process of PBT and ITR in Fig. 3. ITR, due to its computation efficiency, can
 224 afford to run longer iterations and tolerate larger exploration noises. Hence, it can converge easily to
 225 diverse solutions by imposing a large diversity constraint. PBT, however, only converges when the
 226 exploration is faint, otherwise it diverges or converges too slowly.

Table 2: The number of discovered landmarks across 6 seeds with standard deviation in the bracket.

| setting | PBT | ITR |
|-----------|-----------|------------------|
| $N_L = 4$ | 2.0 (1.0) | 3.5 (0.5) |
| $N_L = 5$ | 2.2 (0.9) | 4.5 (0.5) |

227 4.3 Practical Remark

228 Based on the above analyses, we suggest ITR and diversity measures based on state distances
 229 be *preferred* in RL applications. We also acknowledge that, by the no-free-lunch theorem, they
 230 cannot be universal solutions and that trade-offs may still exist (see discussions in Sec.7 and App.F).
 231 Nonetheless, in the following sections, we will show that the effective implementation of these choices
 232 can lead to superior performances in various challenging benchmarks. We hope that our approach will
 233 serve as a starting point and provide valuable insights into the development of increasingly powerful
 234 algorithms for potentially more challenging scenarios.

235 5 Method

236 In this section, we develop a diversity-driven RL algorithm, *State-based Intrinsic-reward Policy*
 237 *Optimization (SIPO)*, by combining ITR and state-distance-based measures. SIPO runs M iterations
 238 to discover M distinct policies. At the i -th iteration, we solve equation (4) by converting it into
 239 unconstrained optimization using the Lagrange method. The unconstrained optimization can be
 240 written as:

$$\min_{\pi_i} \max_{\lambda_j \geq 0, 1 \leq j < i} -J(\pi_i) - \sum_{j=1}^{i-1} \lambda_j (D_S(\pi_i, \pi_j^*) - \delta) \quad (7)$$

241 λ_j ($1 \leq j < i$) are Lagrange multipliers. $\{\pi_j^*\}_{j=1}^{i-1}$ are previously obtained policies. We adopt
 242 two-timescale Gradient Descent Ascent (GDA) [27] to solve the above minimax optimization, i.e.,
 243 performing gradient descent over π_i and gradient ascent over λ_j with different learning rates. In our
 244 algorithm, we additionally enforce the dual variables λ_j to be bounded (i.e., in an interval $[0, \Lambda]$ for
 245 a large number Λ), which plays an important role both in the theoretical analysis and in empirical
 246 convergence. However, $D_S(\pi_i, \pi_j^*)$ cannot be directly optimized w.r.t. π_i through gradient-based
 247 methods because it is related to the states visited by π_i . We cast $D_S(\pi_i, \pi_j^*)$ as intrinsic rewards and
 248 optimize the joint return via policy gradient. The pseudocode of SIPO can be found in App. G.

250 An important property of SIPO is the convergence guarantee. We present an informal illustration in
 251 Thm. 5.1 and present the formal theorem with proof in App. E.2.

252 **Theorem 5.1.** (Informal) *Under continuity assumptions, SIPO converges to an ϵ -stationary point.*

253 **Remark:** We assumed that the return J and the distance D_S are smooth in policies. In practice, this
 254 is true if (1) policy and state space are bounded and (2) reward function and system dynamics are
 255 continuous in the policy. (Continuous functions are bounded over compact spaces.) The key step is to
 256 analyze the role of the bounded dual variables λ , which achieves an $\frac{1}{\Lambda}$ -approximation of constraint
 257 without hurting the optimality condition.

258 Instead of directly defining D_S , we define intrinsic rewards as illustrated in Sec. 5, such that
 259 $D_S(\pi_i, \pi_j^*) = \mathbb{E}_{s_h \sim \mu_{\pi_i}} \left[\sum_{h=1}^H r_{\text{int}}(s_h; \pi_i, \pi_j^*) \right]$.

260 **RBF Kernel:** The most popular realization of Eq. (5) in machine learning is through kernel functions.
 261 Herein, we realize Eq. (5) as an RBF kernel on states. Formally, the intrinsic reward is defined by

$$r_{\text{int}}^{\text{RBF}}(s_h; \pi_i, \pi_j^*) = \frac{1}{H} \mathbb{E}_{s' \sim \mu_{\pi_j^*}} \left[-\exp \left(-\frac{\|s_h - s'\|^2}{2\sigma^2} \right) \right] \quad (8)$$

262 where σ is a hyperparameter controlling the variance.

263 **Wasserstein Distance:** For stronger discrimination power, we can also realize Eq. (5) as L_2 -
 264 Wasserstein distance. According to the dual form [58], we define

$$r_{\text{int}}^{\text{WD}}(s_h; \pi_i, \pi_j^*) = \frac{1}{H} \sup_{\|f\|_L \leq 1} f(s_h) - \mathbb{E}_{s' \sim \mu_{\pi_j^*}} [f(s')] \quad (9)$$

265 where $f : \mathcal{S} \rightarrow \mathbb{R}$ is a 1-Lipschitz function. We implement f as a neural network and clip parameters
 266 to $[-0.01, 0.01]$ to ensure the Lipschitz constraint. Note that $r_{\text{int}}^{\text{WD}}$ incorporates representation learning
 267 by utilizing a learnable scoring function f and is more flexible in practice. We also show in App. B.3
 268 that $r_{\text{int}}^{\text{WD}}$ is robust to different inputs, including states with random noises and RGB images.

269 We name SIPO with $r_{\text{int}}^{\text{RBF}}$ and $r_{\text{int}}^{\text{WD}}$ *SIPO-RBF* and *SIPO-WD* respectively.

270 **Implementation:** To incorporate temporal information, we stack the recent 4 global states to
 271 compute intrinsic rewards and normalize the intrinsic rewards to stabilize training. In multi-agent
 272 environments, we learn an agent-ID-conditioned policy [15] and share the parameter across all agents.
 273 Our implementation is based on MAPPO [64] with more details in App. D.

274 6 Experiments

275 We evaluate SIPO across three domains that exhibit multi-modality of solutions. The first domain
 276 is the humanoid locomotion task in Isaac Gym [38], where diversity can be quantitatively assessed
 277 by well-defined behavior descriptors. We remark that the issues we addressed in Sec. 4.1 may not
 278 be present in this task where the action space is small and actions are highly correlated with states.
 279 Further, we examine the effectiveness of SIPO in two much more challenging multi-agent domains,
 280 StarCarft Multi-Agent Challenge (SMAC) [53] and Google Research Football (GRF) [23], where
 281 well-defined behavior descriptors are not available and existing diversity measures may produce
 282 misleading diversity scores. We provide introductions to these environments in App. C.

283 First, we show that SIPO can efficiently learn diverse strategies and outperform several baseline
 284 methods, including DIPG [39], SMERL [22], DvD [48], and RSPO [69]. Then, we qualitatively
 285 demonstrate the emergent behaviors learned by SIPO, which are both *visually distinguishable* and
 286 *human-interpretable*. Finally, we perform an ablation study over the building components of SIPO
 287 and show that both the diversity measure, ITR, and GDA are critical to the performance.

288 All algorithms run for the same number of environment frames on a desktop machine with an
 289 RTX3090 GPU. Numbers are average values over 5 seeds in Humanoid and SMAC and 3 seeds in
 290 GRF with standard deviation shown in brackets. More algorithm details can be found in App. D.
 291 Additional visualization results can be found in our project website (see App. A).

292 6.1 Comparison with Baseline Methods

293 **Humanoid Locomotion.** Following Zhou et al. [69], we train
 294 a population of size 4. We assess diversity by the pairwise
 295 distance of joint torques, a widely used behavior descriptor
 296 in recent Quality-Diversity works [62]. Torque states are not
 297 included as the input of diversity measures and we only use
 298 them for evaluation to ensure a fair comparison. Results are
 299 shown in Table 3. We can see that both variants of SIPO
 300 can outperform all baseline methods except that SIPO-RBF
 301 achieves comparable performance with RSPO, even if RSPO
 302 explicitly encourages the output of different actions/forces.

Table 3: Pairwise distance of joint torques (i.e., diversity scores) in the humanoid locomotion task.

| SIPO-RBF | SIPO-WD | RSPO |
|------------|-------------------|------------|
| 0.53(0.17) | 0.71(0.23) | 0.53(0.05) |
| DIPG | DvD | SMERL |
| 0.12(0.04) | 0.40(0.22) | 0.01(0.00) |

303 **SMAC** Following Zhou et al. [69], we run SIPO and all baselines on an easy map, *2m_vs_1z*, and
 304 a hard map, *2c_vs_64zg*, both across 4 iterations. We merge all trajectories produced by the policy
 305 collection and incorporate a k -nearest-neighbor state entropy estimation [28] to assess diversity.
 306 Intuitively, a more diverse population should have a larger state entropy value. We set $k = 12$
 307 following Liu and Abbeel [28] and show results in Table 4. On these maps, two agents are both
 308 involved in the attack. Therefore, RSPO, which incorporates an action-based cross-entropy measure,
 309 can perform well across all baselines. However, SIPO explicitly compares the distance between
 310 resulting trajectories and can even outperform RSPO, leading to the most diverse population.

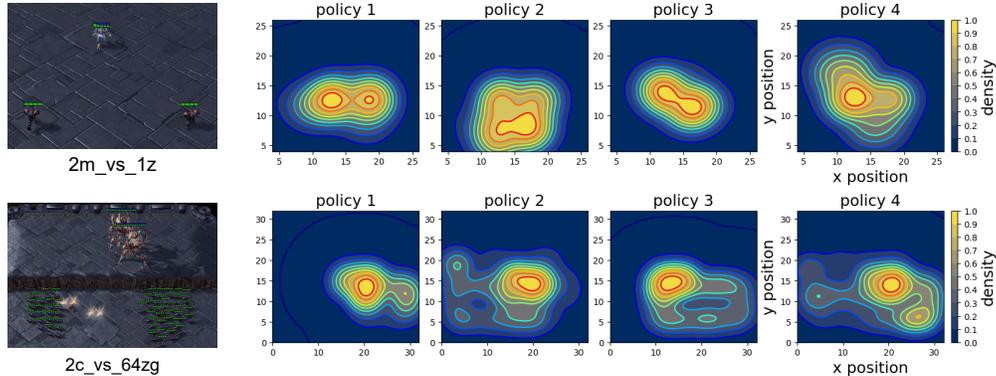


Figure 5: Heatmaps of agent positions in SMAC across 4 iterations with SIPO-RBF.

311 **GRF** We consider three academy scenarios, specifically *3v1*, *counterattack (CA)*, and *corner*. The
 312 GRF environment is more challenging than SMAC due to the large action space, more agents, and the
 313 existence of duplicate actions. We determine a population size $M = 4$ by balancing resources and
 314 wall-clock time across different baselines. Table 5 compares the number of distinct policies (in terms
 315 of ball-passing routes, see App. B.2) discovered in the population. Due to the strong adversarial power
 316 of our diversity measures and the application of GDA, SIPO is the most efficient and robust — even
 317 in the challenging 11-vs-11 *corner* and *CA* scenario, SIPO can effectively discover different winning
 318 strategies in just a few iterations across different seeds. By contrast, baselines suffer from learning
 319 instability in these challenging environments and tend to discover policies with slight distinctions.
 320 We also calculate the estimated state entropy as we did in SMAC. However, we find that this metric
 321 cannot distinguish fine-grained ball-passing behaviors in GRF (check our discussions in App. B).

322 **Remark:** In GRF experiments, when M is small, even
 323 repeated training with different random seeds (PG) is a
 324 strong baseline (see Table 5). Hence, the numbers are
 325 actually restricted in a small interval (with a lower bound
 326 equal to PG results and an upper bound equal to $M = 4$),
 327 which makes the improvements by SIPO seemingly less
 328 significant. However, achieving clear improvements in
 329 these challenging applications remains particularly non-
 330 trivial. With a population size $M = 10$, SIPO clearly
 331 outperforms baselines by consistently discovering one or
 332 more additional strategies.

Table 4: State entropy estimated by k -nearest-neighbor in SMAC. ($k = 12$)

| | <i>2m_vs_1z</i> | <i>2c_vs_64zg</i> |
|----------|---------------------|---------------------|
| SIPO-RBF | 0.038(0.002) | 0.072(0.003) |
| SIPO-WD | 0.036(0.001) | 0.056(0.003) |
| RSPO | 0.032(0.003) | 0.070(0.001) |
| DIPG | 0.032(0.002) | 0.056(0.004) |
| SMERL | 0.028(0.002) | 0.042(0.002) |
| DvD | 0.030(0.002) | 0.057(0.003) |

333 6.2 Qualitative Analysis

334 For SMAC, we present heatmaps of agent positions in Fig. 5. The heatmaps clearly show that SIPO
 335 can consistently learn novel winning strategies to conquer the enemy. Fig. 6 presents the learned
 336 behavior by SIPO in the GRF *3v1* scenario of seed 1. We can observe that agents have learned a wide
 337 spectrum of collaboration strategies across merely 7 iterations. The strategies discovered by SIPO
 338 are both *diverse* and *human-interpretable*. In the first iteration, all agents are involved in the attack
 339 such that they can distract the defender and obtain a high win rate. The 2nd and the 6th iteration
 340 demonstrate an efficient pass-and-shoot strategy, where agents quickly elude the defender and score a
 341 goal. In the 3rd and the 7th iterations, agents learn smart “one-two” strategies to bypass the defender,
 342 a prevalent tactic employed by human football players. We note that *NONE* of the baselines have

Table 5: Number of distinct strategies in GRF discovered by different methods in terms of the ball-passing route. Details of the evaluation protocol can be found in App. B.2.

| | Population Size M | ours | | baselines | | | | random |
|---------------|---------------------|------------------|------------------|-----------|-----------|------------------|-----------|-----------|
| | | SIPO-RBF | SIPO-WD | DIPG | SMERL | DvD ¹ | RSPO | PG |
| <i>3v1</i> | 4 | 3.0 (0.8) | 3.0 (0.0) | 2.7 (0.5) | 1.3 (0.5) | 3.0 (0.8) | 2.0 (0.0) | 2.7 (0.5) |
| <i>CA</i> | 4 | 3.3 (0.5) | 3.0 (0.8) | 2.3 (0.5) | 1.3 (0.5) | - | 2.0 (0.0) | 1.7 (0.5) |
| <i>corner</i> | 4 | 2.7 (0.5) | 3.0 (0.8) | 1.7 (0.5) | 1.0 (0.0) | - | 1.6 (0.5) | 2.0 (0.8) |
| <i>3v1</i> | 10 | 4.3 (0.5) | 5.7 (0.5) | 3.7 (0.5) | - | - | 2.3 (0.5) | - |

¹ Training DvD in *CA* and *corner* or with $M = 10$ requires >24GB GPU memory, which exceeds our memory limit.

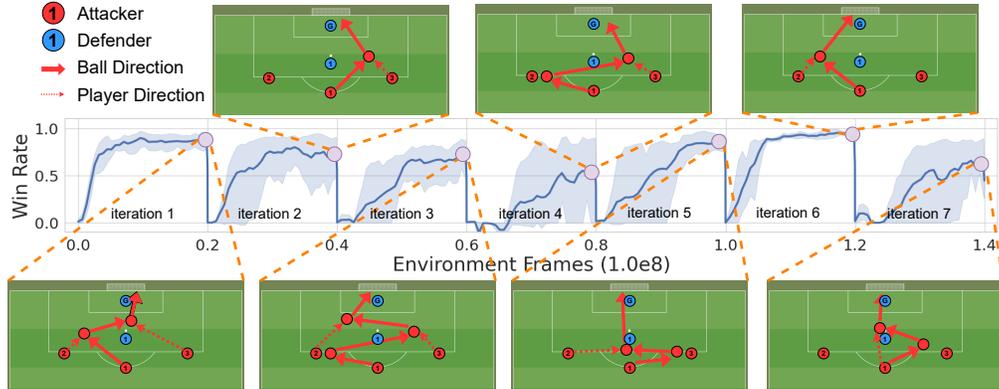


Figure 6: Learning curves and discovered strategies by SIPO-WD in the *3v1* scenario over 7 iterations. Strategies of seed 1 are shown.

343 ever discovered this strategy across all runs, while SIPO is consistently able to derive such strategies
 344 for all random seeds. Visualization results in *CA* and *corner* scenarios can be found in App. B.

345 6.3 Ablation Study

346 We apply these changes to SIPO-WD:

- 347 • *fix-L*: Fixing the multiplier λ_i instead of applying GDA.
- 348 • *CE*: The intrinsic reward is replaced with cross-entropy, i.e., $r_{\text{int}}^{\text{CE}}(s_h, a_h) = -\log \pi_j^*(a_h | s_h)$, where
 349 π_j^* denotes a previously discovered policy. Additionally, GDA is still applied.
- 350 • *filter*: Optimizing the extrinsic rewards on trajectories that have intrinsic returns exceeding δ and
 351 optimizing intrinsic rewards defined by Eq. (9) for other trajectories [69].
- 352 • *PBT*: Simultaneously training M policies with $M(M-1)/2$ constraints (i.e., directly solving Eq. (3))
 353 with intrinsic rewards defined by Eq. (9) and GDA.

354 We report the number of visually
 355 distinct policies discovered by these
 356 methods in Table 6. Comparison be-
 357 tween SIPO and CE demonstrates
 358 that the action-based cross-entropy
 359 measure may suffer from duplicate
 360 actions in GRF and produce nearly
 361 identical behavior by overly exploit-
 362 ing duplicate actions, especially in the *CA* and *corner* scenarios with 11 agents. Besides, the fixed
 363 Lagrange coefficient, the filtering-based method, and PBT are all detrimental to our algorithm. These
 364 methods also suffer from significant training instability. Overall, the state-distance-based diversity
 365 measure, ITR, and GDA are all critical to the performance of SIPO.

Table 6: # distinct strategies of ablations in GRF.

| | ours | fix-L | CE | filter | PBT |
|---------------|------------------|--------------|-----------|-----------|--------------|
| <i>3v1</i> | 3.0 (0.0) | 1.0 (0.0) | 2.7 (0.5) | 1.3 (0.5) | 2.7 (0.5) |
| <i>CA</i> | 3.0 (0.8) | ¹ | 2.3 (0.8) | 1.0 (0.0) | ² |
| <i>corner</i> | 3.0 (0.8) | ¹ | 1.7 (0.5) | 1.0 (0.0) | ² |

¹ Not converged.

² Training requires >24GB memory and exceeds our memory limit.

366 7 Conclusion

367 We tackle the problem of discovering diverse high-reward policies in RL. First, we demonstrate
 368 concrete failure cases of existing diversity measures and propose a novel measure that explicitly
 369 compares the distance in state space. Next, we present a thorough comparison between PBT and ITR,
 370 and show that ITR is much easier to optimize and can derive solutions with comparable quality to
 371 PBT. Motivated by these insights, we combine ITR with a state-distance-based diversity measure
 372 to develop SIPO, which has provable convergence and can efficiently discover a wide spectrum of
 373 human-interpretable strategies in a wide range of environments.

374 **Limitations:** First, we assume direct access to an object-centric state representation. When such
 375 a representation is not available (e.g., image-based observations), representation learning becomes
 376 necessary and algorithm performance can be affected by the quality of the learned representations.
 377 Second, because ITR requires sequential training, the wall clock time of SIPO can be longer than the
 378 PBT alternatives when fixing the total number of training samples. The acceleration of ITR remains
 379 an open challenge.

380 References

- 381 [1] Monica Babes, Enrique Munoz de Cote, and Michael L. Littman. Social reward shaping in
382 the prisoner’s dilemma. In Lin Padgham, David C. Parkes, Jörg P. Müller, and Simon Parsons,
383 editors, *7th International Joint Conference on Autonomous Agents and Multiagent Systems*
384 (*AAMAS 2008*), Estoril, Portugal, May 12-16, 2008, Volume 3, pages 1389–1392. IFAAMAS,
385 2008. URL <https://dl.acm.org/citation.cfm?id=1402880>.
- 386 [2] Bowen Baker, Ingmar Kanitscheider, Todor M. Markov, Yi Wu, Glenn Powell, Bob McGrew,
387 and Igor Mordatch. Emergent tool use from multi-agent autocurricula. In *8th International*
388 *Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.
389 OpenReview.net, 2020. URL <https://openreview.net/forum?id=SkxpxJBkWS>.
- 390 [3] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi
391 Munos. Unifying count-based exploration and intrinsic motivation. *Advances in neural*
392 *information processing systems*, 29, 2016.
- 393 [4] Yura Burda and Harri Edwards, Oct 2018. URL [https://openai.com/blog/
394 reinforcement-learning-with-prediction-based-rewards/](https://openai.com/blog/reinforcement-learning-with-prediction-based-rewards/).
- 395 [5] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random
396 network distillation. *arXiv preprint arXiv:1810.12894*, 2018.
- 397 [6] Andres Campero, Roberta Raileanu, Heinrich Küttler, Joshua B. Tenenbaum, Tim Rocktäschel,
398 and Edward Grefenstette. Learning with amigo: Adversarially motivated intrinsic goals. In
399 *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria,*
400 *May 3-7, 2021*. OpenReview.net, 2021. URL [https://openreview.net/forum?id=ETBc_
401 MIMgoX](https://openreview.net/forum?id=ETBc_MIMgoX).
- 402 [7] Víctor Campos, Alexander Trott, Caiming Xiong, Richard Socher, Xavier Giró-i Nieto, and
403 Jordi Torres. Explore, discover and learn: Unsupervised discovery of state-covering skills. In
404 *International Conference on Machine Learning*, pages 1317–1327. PMLR, 2020.
- 405 [8] Rujikorn Charakorn, Poramate Manoonpong, and Nat Dilokthanakul. Generating diverse
406 cooperative agents by learning incompatible policies. In *ICML 2022 Workshop AI for Agent-*
407 *Based Modelling*, 2022.
- 408 [9] Jack Clark and Dario Amodei, Dec 2016. URL [https://openai.com/blog/
409 faulty-reward-functions/](https://openai.com/blog/faulty-reward-functions/).
- 410 [10] Brandon Cui, Andrei Lupu, Samuel Sokota, Hengyuan Hu, David J Wu, and Jakob Nicolaus
411 Foerster. Adversarial diversity in hanabi. In *The Eleventh International Conference on Learning*
412 *Representations*, 2023.
- 413 [11] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt
414 like animals. *Nature*, 521(7553):503–507, May 2015. ISSN 0028-0836, 1476-4687. doi:
415 10.1038/nature14422. URL <http://www.nature.com/articles/nature14422>.
- 416 [12] Kalyanmoy Deb and Amit Saha. Finding multiple solutions for multimodal optimization prob-
417 lems using a multi-objective evolutionary approach. In Martin Pelikan and Jürgen Branke, edi-
418 tors, *Genetic and Evolutionary Computation Conference, GECCO 2010, Proceedings, Portland,*
419 *Oregon, USA, July 7-11, 2010*, pages 447–454. ACM, 2010. doi: 10.1145/1830483.1830568.
420 URL <https://doi.org/10.1145/1830483.1830568>.
- 421 [13] Sam Devlin and Daniel Kudenko. Theoretical considerations of potential-based reward shaping
422 for multi-agent systems. In Liz Sonenberg, Peter Stone, Kagan Tumer, and Pinar Yolum, editors,
423 *10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011),*
424 *Taipei, Taiwan, May 2-6, 2011, Volume 1-3*, pages 225–232. IFAAMAS, 2011. URL [http:
425 //portal.acm.org/citation.cfm?id=2030503&CFID=69153967&CFTOKEN=38069692](http://portal.acm.org/citation.cfm?id=2030503&CFID=69153967&CFTOKEN=38069692).
- 426 [14] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you
427 need: Learning skills without a reward function. In *7th International Conference on Learning*
428 *Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
429 URL <https://openreview.net/forum?id=SJx63jRqFm>.

- 430 [15] Wei Fu, Chao Yu, Zelai Xu, Jiaqi Yang, and Yi Wu. Revisiting some common prac-
431 tices in cooperative multi-agent reinforcement learning. In Kamalika Chaudhuri, Stefanie
432 Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings*
433 *of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings*
434 *of Machine Learning Research*, pages 6863–6877. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/fu22d.html>.
435
- 436 [16] Dibya Ghosh, Abhishek Gupta, and Sergey Levine. Learning actionable representations with
437 goal conditioned policies. In *7th International Conference on Learning Representations,*
438 *ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=Hye91nCct7>.
439
- 440 [17] Agrim Gupta, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. Embodied intelligence via
441 learning and evolution. *Nature communications*, 12(1):1–12, 2021.
- 442 [18] Elad Hazan, Sham Kakade, Karan Singh, and Abby Van Soest. Provably efficient maximum
443 entropy exploration. In *International Conference on Machine Learning*, pages 2681–2691.
444 PMLR, 2019.
- 445 [19] Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M. Czarnecki, Jeff Donahue, Ali
446 Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, Chrisantha Fernando, and
447 Koray Kavukcuoglu. Population Based Training of Neural Networks, November 2017. URL
448 <http://arxiv.org/abs/1711.09846>. arXiv:1711.09846 [cs].
- 449 [20] Max Jaderberg, Wojciech M. Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia
450 Castañeda, Charles Beattie, Neil C. Rabinowitz, Ari S. Morcos, Avraham Ruderman, Nicolas
451 Sonnerat, Tim Green, Louise Deason, Joel Z. Leibo, David Silver, Demis Hassabis, Koray
452 Kavukcuoglu, and Thore Graepel. Human-level performance in 3D multiplayer games with
453 population-based reinforcement learning. *Science*, 364(6443):859–865, May 2019. ISSN
454 0036-8075, 1095-9203. doi: 10.1126/science.aau6249. URL <https://www.science.org/doi/10.1126/science.aau6249>.
455
- 456 [21] Zheyuan Jiang, Jingyue Gao, and Jianyu Chen. Unsupervised skill discovery via recurrent skill
457 training. *Advances in Neural Information Processing Systems*, 35:39034–39046, 2022.
- 458 [22] Saurabh Kumar, Aviral Kumar, Sergey Levine, and Chelsea Finn. One solution
459 is not all you need: Few-shot extrapolation via structured maxent RL. In Hugo
460 Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien
461 Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Confer-*
462 *ence on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-*
463 *12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/5d151d1059a6281335a10732fc49620e-Abstract.html>.
464
- 465 [23] Karol Kurach, Anton Raichuk, Piotr Stanczyk, Michal Zajac, Olivier Bachem, Lasse Espeholt,
466 Carlos Riquelme, Damien Vincent, Marcin Michalski, Olivier Bousquet, and Sylvain Gelly.
467 Google research football: A novel reinforcement learning environment. In *The Thirty-Fourth*
468 *AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Ap-*
469 *plications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on*
470 *Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12,*
471 *2020*, pages 4501–4510. AAAI Press, 2020. URL <https://ojs.aaai.org/index.php/AAAI/article/view/5878>.
472
- 473 [24] Yoonho Lee, Huaxiu Yao, and Chelsea Finn. Diversify and disambiguate: Learning from
474 underspecified data. *CoRR*, abs/2202.03418, 2022. URL <https://arxiv.org/abs/2202.03418>.
475
- 476 [25] Chenghao Li, Tonghan Wang, Chengjie Wu, Qianchuan Zhao, Jun Yang, and Chongjie Zhang.
477 Celebrating diversity in shared multi-agent reinforcement learning. In Marc’Aurelio Ran-
478 zato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan,
479 editors, *Advances in Neural Information Processing Systems 34: Annual Conference on*
480 *Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*,
481 pages 3991–4002, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/20aee3a5f4643755a79ee5f6a73050ac-Abstract.html>.
482

- 483 [26] Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. Deep
484 reinforcement learning for dialogue generation. In Jian Su, Xavier Carreras, and Kevin Duh,
485 editors, *Proceedings of the 2016 Conference on Empirical Methods in Natural Language
486 Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1192–1202. The
487 Association for Computational Linguistics, 2016. doi: 10.18653/v1/d16-1127. URL <https://doi.org/10.18653/v1/d16-1127>.
488
- 489 [27] Tianyi Lin, Chi Jin, and Michael I. Jordan. On gradient descent ascent for nonconvex-concave
490 minimax problems. In *Proceedings of the 37th International Conference on Machine Learning,
491 ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning
492 Research*, pages 6083–6093. PMLR, 2020. URL [http://proceedings.mlr.press/v119/
493 lin20a.html](http://proceedings.mlr.press/v119/lin20a.html).
- 494 [28] Hao Liu and Pieter Abbeel. Behavior from the void: Unsupervised active pre-training. *Advances
495 in Neural Information Processing Systems*, 34:18459–18473, 2021.
- 496 [29] Iou-Jen Liu, Unnat Jain, Raymond A. Yeh, and Alexander G. Schwing. Cooperative explo-
497 ration for multi-agent deep reinforcement learning. In Marina Meila and Tong Zhang, editors,
498 *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24
499 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages
500 6826–6836. PMLR, 2021. URL <http://proceedings.mlr.press/v139/liu21j.html>.
- 501 [30] Siqi Liu, Guy Lever, Josh Merel, Saran Tunyasuvunakool, Nicolas Heess, and Thore Graepel.
502 Emergent coordination through competition. In *7th International Conference on Learning
503 Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
504 URL <https://openreview.net/forum?id=BkG8sJR5Km>.
- 505 [31] Siqi Liu, Guy Lever, Zhe Wang, Josh Merel, S. M. Ali Eslami, Daniel Hennes, Wojciech M.
506 Czarnecki, Yuval Tassa, Shayegan Omidshafiei, Abbas Abdolmaleki, Noah Y. Siegel, Leonard
507 Hasenclever, Luke Marris, Saran Tunyasuvunakool, H. Francis Song, Markus Wulfmeier, Paul
508 Muller, Tuomas Haarnoja, Brendan D. Tracey, Karl Tuyls, Thore Graepel, and Nicolas Heess.
509 From motor control to team play in simulated humanoid football. *Sci. Robotics*, 7(69), 2022. doi:
510 10.1126/scirobotics.abo0235. URL <https://doi.org/10.1126/scirobotics.abo0235>.
- 511 [32] Xiangyu Liu, Hangtian Jia, Ying Wen, Yujing Hu, Yingfeng Chen, Changjie Fan, Zhipeng Hu,
512 and Yaodong Yang. Towards unifying behavioral and response diversity for open-ended learning
513 in zero-sum games. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy
514 Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing
515 Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021,
516 December 6-14, 2021, virtual*, pages 941–952, 2021. URL [https://proceedings.neurips.
517 cc/paper/2021/hash/07bba581a2dd8d098a3be0f683560643-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/07bba581a2dd8d098a3be0f683560643-Abstract.html).
- 518 [33] Qian Long, Zihan Zhou, Abhinav Gupta, Fei Fang, Yi Wu, and Xiaolong Wang. Evolutionary
519 population curriculum for scaling multi-agent reinforcement learning. In *International
520 Conference on Learning Representations*, 2020.
- 521 [34] Andrei Lupu, Hengyuan Hu, and Jakob N. Foerster. Trajectory diversity for zero-shot co-
522 ordination. In Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé, editors,
523 *AAMAS ’21: 20th International Conference on Autonomous Agents and Multiagent Sys-
524 tems, Virtual Event, United Kingdom, May 3-7, 2021*, pages 1593–1595. ACM, 2021. doi:
525 10.5555/3463952.3464170. URL [https://www.ifaamas.org/Proceedings/aamas2021/
526 pdfs/p1593.pdf](https://www.ifaamas.org/Proceedings/aamas2021/pdfs/p1593.pdf).
- 527 [35] Pingchuan Ma, Tao Du, and Wojciech Matusik. Efficient continuous pareto exploration in
528 multi-task learning. In *Proceedings of the 37th International Conference on Machine Learning,
529 ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning
530 Research*, pages 6522–6531. PMLR, 2020. URL [http://proceedings.mlr.press/v119/
531 ma20a.html](http://proceedings.mlr.press/v119/ma20a.html).
- 532 [36] Tengyu Ma. Why Do Local Methods Solve Nonconvex Problems?, March 2021. URL [http://
533 arxiv.org/abs/2103.13462](http://arxiv.org/abs/2103.13462). arXiv:2103.13462 [cs, math, stat].

- 534 [37] Marlos C. Machado, Marc G. Bellemare, and Michael Bowling. Count-based exploration with
535 the successor representation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence,*
536 *AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference,*
537 *IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence,*
538 *EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 5125–5133. AAAI Press, 2020.
539 URL <https://ojs.aaai.org/index.php/AAAI/article/view/5955>.
- 540 [38] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles
541 Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac
542 gym: High performance gpu-based physics simulation for robot learning, 2021.
- 543 [39] Muhammad A. Masood and Finale Doshi-Velez. Diversity-inducing policy gradient: Using
544 maximum mean discrepancy to find a set of diverse policies. In Sarit Kraus, editor, *Proceedings*
545 *of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019,*
546 *Macao, China, August 10-16, 2019*, pages 5923–5929. ijcai.org, 2019. doi: 10.24963/ijcai.
547 2019/821. URL <https://doi.org/10.24963/ijcai.2019/821>.
- 548 [40] B.L. Miller and M.J. Shaw. Genetic algorithms with dynamic niche sharing for multimodal
549 function optimization. In *Proceedings of IEEE International Conference on Evolutionary*
550 *Computation*, pages 786–791, Nagoya, Japan, 1996. IEEE. ISBN 978-0-7803-2902-7. doi:
551 10.1109/ICEC.1996.542701. URL <http://ieeexplore.ieee.org/document/542701/>.
- 552 [41] Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *CoRR*,
553 abs/1504.04909, 2015. URL <http://arxiv.org/abs/1504.04909>.
- 554 [42] Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transfor-
555 mations: Theory and application to reward shaping. In *Icml*, volume 99, pages 278–287,
556 1999.
- 557 [43] Olle Nilsson and Antoine Cully. Policy gradient assisted MAP-Elites. In *Proceedings of the*
558 *Genetic and Evolutionary Computation Conference*, pages 866–875, Lille France, June 2021.
559 ACM. ISBN 978-1-4503-8350-9. doi: 10.1145/3449639.3459304. URL <https://dl.acm.org/doi/10.1145/3449639.3459304>.
- 561 [44] Shayegan Omidshafiei, Karl Tuyls, Wojciech M Czarnecki, Francisco C Santos, Mark Rowland,
562 Jerome Connor, Daniel Hennes, Paul Muller, Julien Pérolat, Bart De Vylder, et al. Navigating
563 the landscape of multiplayer games. *Nature communications*, 11(1):1–17, 2020.
- 564 [45] Takayuki Osa, Voot Tangkaratt, and Masashi Sugiyama. Discovering diverse solutions in deep
565 reinforcement learning by maximizing state-action-based mutual information. *Neural Networks*,
566 152:90–104, 2022. doi: 10.1016/j.neunet.2022.04.009. URL <https://doi.org/10.1016/j.neunet.2022.04.009>.
- 568 [46] Aldo Pacchiano, Jack Parker-Holder, Yunhao Tang, Krzysztof Choromanski, Anna Choro-
569 manska, and Michael Jordan. Learning to score behaviors for guided policy optimization. In
570 *International Conference on Machine Learning*, pages 7445–7454. PMLR, 2020.
- 571 [47] Jack Parker-Holder, Luke Metz, Cinjon Resnick, Hengyuan Hu, Adam Lerer, Alistair Letcher,
572 Alexander Peysakhovich, Aldo Pacchiano, and Jakob N. Foerster. Ridge rider: Finding diverse
573 solutions by following eigenvectors of the hessian. In Hugo Larochelle, Marc’Aurelio Ranzato,
574 Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Informa-*
575 *tion Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020,*
576 *NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/08425b881bcde94a383cd258cea331be-Abstract.html>.
- 578 [48] Jack Parker-Holder, Aldo Pacchiano, Krzysztof Marcin Choromanski, and Stephen J.
579 Roberts. Effective diversity in population based reinforcement learning. In Hugo
580 Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien
581 Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Confer-*
582 *ence on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-*
583 *12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/d1dc3a8270a6f9394f88847d7f0050cf-Abstract.html>.

- 585 [49] Alexander Peysakhovich and Adam Lerer. Consequentialist conditional cooperation in social
586 dilemmas with imperfect information. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=BkabRiQpb>.
- 589 [50] Thomas Pierrot, Valentin Macé, Félix Chalumeau, Arthur Flajolet, Geoffrey Cideron, Karim
590 Beguir, Antoine Cully, Olivier Sigaud, and Nicolas Perrin-Gilbert. Diversity Policy Gradient
591 for Sample Efficient Quality-Diversity Optimization. In *Proceedings of the Genetic and
592 Evolutionary Computation Conference*, pages 1075–1083, July 2022. doi: 10.1145/3512290.
593 3528845. URL <http://arxiv.org/abs/2006.08505>. arXiv:2006.08505 [cs].
- 594 [51] Justin K. Pugh, Lisa B. Soros, and Kenneth O. Stanley. Quality Diversity: A New Frontier for
595 Evolutionary Computation. *Frontiers in Robotics and AI*, 3, July 2016. ISSN 2296-9144. doi:
596 10.3389/frobt.2016.00040. URL [http://journal.frontiersin.org/Article/10.3389/
597 frobt.2016.00040/abstract](http://journal.frontiersin.org/Article/10.3389/frobt.2016.00040/abstract).
- 598 [52] Tim Roughgarden, editor. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University
599 Press, 2020. ISBN 9781108637435. doi: 10.1017/9781108637435. URL [https://doi.org/
600 10.1017/9781108637435](https://doi.org/10.1017/9781108637435).
- 601 [53] Mikayel Samvelyan, Tabish Rashid, Christian Schröder de Witt, Gregory Farquhar, Nantas
602 Nardelli, Tim G. J. Rudner, Chia-Man Hung, Philip H. S. Torr, Jakob N. Foerster, and Shi-
603 mon Whiteson. The starcraft multi-agent challenge. In Edith Elkind, Manuela Veloso, Noa
604 Agmon, and Matthew E. Taylor, editors, *Proceedings of the 18th International Conference on
605 Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17,
606 2019*, pages 2186–2188. International Foundation for Autonomous Agents and Multiagent
607 Systems, 2019. URL <http://dl.acm.org/citation.cfm?id=3332052>.
- 608 [54] Matthijs TJ Spaan. Partially observable markov decision processes. In *Reinforcement Learning*,
609 pages 387–414. Springer, 2012.
- 610 [55] Hao Sun, Zhenghao Peng, Bo Dai, Jian Guo, Dahua Lin, and Bolei Zhou. Novel policy seeking
611 with constrained optimization. *arXiv preprint arXiv:2005.10696*, 2020.
- 612 [56] Zhenggang Tang, Chao Yu, Boyuan Chen, Huazhe Xu, Xiaolong Wang, Fei Fang, Simon Shaolei
613 Du, Yu Wang, and Yi Wu. Discovering diverse multi-agent strategic behavior via reward
614 randomization. In *9th International Conference on Learning Representations, ICLR 2021,
615 Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL [https://openreview.
616 net/forum?id=lvRTC669EY_](https://openreview.net/forum?id=lvRTC669EY_).
- 617 [57] Luca Venturi, Afonso S. Bandeira, and Joan Bruna. Neural networks with finite intrinsic
618 dimension have no spurious valleys. *CoRR*, abs/1802.06384, 2018. URL [http://arxiv.org/
619 abs/1802.06384](http://arxiv.org/abs/1802.06384).
- 620 [58] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.
- 621 [59] Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik,
622 Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh,
623 Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P.
624 Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dal-
625 ibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Ziyu Wang,
626 Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney,
627 Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps,
628 and David Silver. Grandmaster level in StarCraft II using multi-agent reinforcement learning.
629 *Nature*, 575(7782):350–354, November 2019. ISSN 0028-0836, 1476-4687. doi: 10.1038/
630 s41586-019-1724-z. URL <http://www.nature.com/articles/s41586-019-1724-z>.
- 631 [60] Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. Poet: open-ended coevolution of
632 environments and their optimized solutions. In *Proceedings of the Genetic and Evolutionary
633 Computation Conference*, pages 142–151, 2019.

- 634 [61] Tonghan Wang, Tarun Gupta, Anuj Mahajan, Bei Peng, Shimon Whiteson, and Chongjie Zhang.
635 RODE: learning roles to decompose multi-agent tasks. In *9th International Conference on*
636 *Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net,
637 2021. URL <https://openreview.net/forum?id=TTUVg6vkNjK>.
- 638 [62] Shuang Wu, Jian Yao, Haobo Fu, Ye Tian, Chao Qian, Yaodong Yang, QIANG FU, and Yang
639 Wei. Quality-similar diversity via population based reinforcement learning. In *The Eleventh*
640 *International Conference on Learning Representations*, 2023.
- 641 [63] Yifan Wu, George Tucker, and Ofir Nachum. The laplacian in RL: learning representations
642 with efficient approximations. In *7th International Conference on Learning Representations,*
643 *ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=HJ1NpoA5YQ>.
- 645 [64] Chao Yu, Akash Velu, Eugene Vinitisky, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising
646 effectiveness of ppo in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*, 2021.
- 647 [65] Tom Zahavy, Brendan O’Donoghue, Andre Barreto, Volodymyr Mnih, Sebastian Flennerhag,
648 and Satinder Singh. Discovering diverse nearly optimal policies with successor features. *arXiv*
649 *preprint arXiv:2106.00669*, 2021.
- 650 [66] Tom Zahavy, Yannick Schroecker, Feryal M. P. Behbahani, Kate Baumli, Sebastian Flennerhag,
651 Shaobo Hou, and Satinder Singh. Discovering policies with domino: Diversity optimization
652 maintaining near optimality. *CoRR*, abs/2205.13521, 2022. doi: 10.48550/arXiv.2205.13521.
653 URL <https://doi.org/10.48550/arXiv.2205.13521>.
- 654 [67] Yunbo Zhang, Wenhao Yu, and Greg Turk. Learning novel policies for tasks. In *International*
655 *Conference on Machine Learning*, pages 7483–7492. PMLR, 2019.
- 656 [68] Rui Zhao, Jinming Song, Hu Haifeng, Yang Gao, Yi Wu, Zhongqian Sun, and Yang Wei.
657 Maximum entropy population based training for zero-shot human-ai coordination. *arXiv*
658 *preprint arXiv:2112.11701*, 2021.
- 659 [69] Zihan Zhou, Wei Fu, Bingliang Zhang, and Yi Wu. Continuously discovering novel strategies
660 via reward-switching policy optimization. In *The Tenth International Conference on Learning*
661 *Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL
662 https://openreview.net/forum?id=hcQHRHKfN_.

663 **A Project Website**

664 Check <https://sites.google.com/view/diversity-sipo> for GIF demonstrations.

665 **B Additional Results**

666 **B.1 More Qualitative Results**

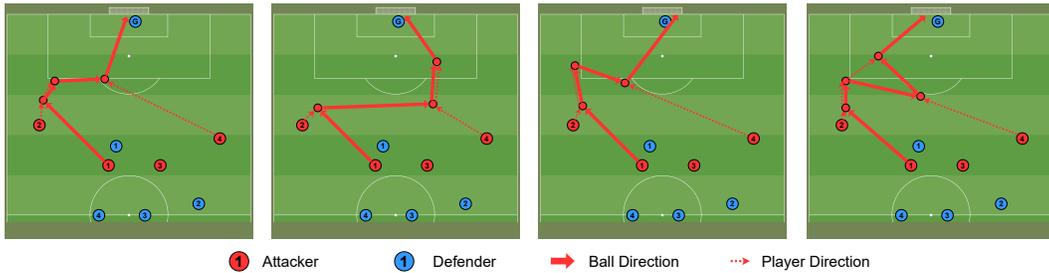


Figure 7: Visualization of learned behaviors in GRF *CA* across a single training trial.

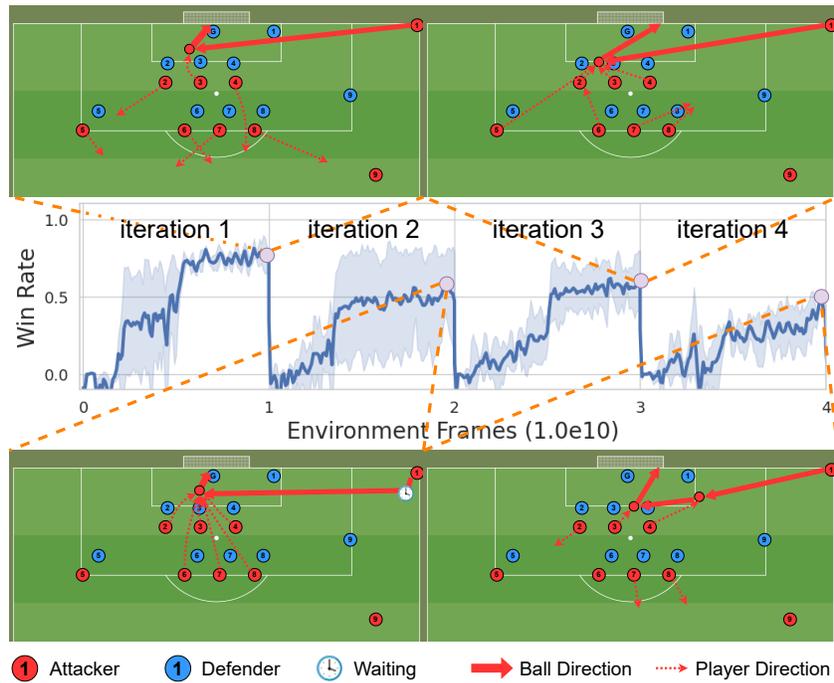


Figure 8: Visualization of learned behaviors in GRF *corner*.

667 We show additional visualization results in Fig. 7, Fig. 8, and Fig. 9. Corresponding GIF visualizations
668 can be found in our project website.

669 **B.2 Evaluation Metric and Protocol**

670 **B.2.1 Humanoid**

671 The Humanoid locomotion task is well-studied in the Quality-Diversity (QD) community, enabling
672 the application of well-defined behavior descriptors (BD) to assess diversity scores. While domain-
673 agnostic metrics like DvD scores can also be applied, we consider domain-specific BDs to be more
674 appropriate and accurate for evaluation in this setting.

Table 7: k -nearest neighbor state entropy estimation in GRF. Population size $M = 4$.

| | ours | | baselines | | | | PG (random seeding) |
|---------------|--------------|--------------|--------------|--------------------|------------------|-------------------|---------------------|
| | SIPO-RBF | SIPO-WD | DIPG | SMERL ¹ | DvD ² | RSPO ¹ | |
| <i>3v1</i> | 0.009(0.000) | 0.012(0.000) | 0.010(0.001) | 0.011(0.002) | 0.010(0.000) | 0.011(0.001) | 0.009(0.001) |
| <i>CA</i> | 0.037(0.000) | 0.031(0.006) | 0.036(0.002) | - | - | 0.034(0.001) | 0.039(0.001) |
| <i>Corner</i> | 0.028(0.001) | 0.031(0.001) | 0.030(0.002) | - | - | - | 0.028(0.002) |

¹The learned policy in some iteration cannot even collect a single winning trajectory. We are unable to compute diversity score.

²Training DvD in *CA* and *corner* requires >24GB GPU memory, which exceeds our memory limit.

675 B.2.2 SMAC

676 Complex multi-agent tasks like SMAC lack well-defined BDs. Hence, domain-agnostic diversity
 677 measures such as the state-entropy measure should be applied. Moreover, different SMAC winning
 678 strategies tend to visit different areas of the map, which can be usually captured by the state-entropy
 679 measure.

680 B.2.3 GRF

681 In our initial study of the GRF task, diversity was evaluated using the k -nearest-neighbor state entropy
 682 estimation as in SMAC (see Table 7). However, we observed a significant difference between the
 683 computed scores and visualized behaviors. Further investigation revealed that state entropy can
 684 sometimes report fake diversity in GRF. For example, the ball-moving route is highly fine-grained
 685 between nearby players in the counter-attack (CA) scenario, and additional passes may not change
 686 the state entropy significantly. Instead, agents’ positions play a crucial role in this scenario, where
 687 different shooting positions can introduce substantial state variance and lead to a higher entropy score.
 688 As an example, readers can refer to the replays of SIPO-RBF (4 iterations of seed 2) and PG (seed 2,
 689 1002, 2002, and 3002), where SIPO-RBF discovers four distinct passing strategies, while PG keeps
 690 passing the ball to the same player. Nevertheless, the state entropy of PG (0.0397) is higher than that
 691 of SIPO-RBF (0.0378).

692 Hence, we counted the number of distinct policies according to their ball-passing routes, such as
 693 passing the ball to different players or shooting with different players, to evaluate diversity in GRF.
 694 To quantify these differences, we extracted the positions of the ball and the players in the field and
 695 calculated the nearest ally player ID to the ball across a winning episode. We then removed timesteps
 696 where the nearest distance was above a pre-defined threshold of 0.03. Typically, these timesteps
 697 correspond to instances when the ball is being transferred among players, making the nearest player
 698 ID irrelevant. Next, we removed consecutive duplicate player IDs from the resulting sequence to
 699 obtain a concise and informative embedding of the ball passing route. By comparing the lengths of
 700 their respective embeddings and verifying that the player IDs in each embedding are identical, we
 701 determined whether two policies exhibit similar behavior.

702 We acknowledge that existing diversity measures may not be applicable in GRF, and hence we
 703 opted for this novel approach to evaluate diversity. Additionally, we experimented with using raw

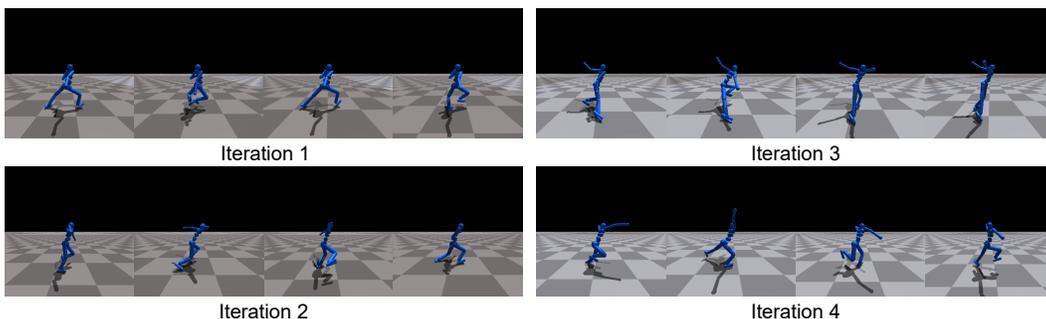


Figure 9: Visualization of learned behaviors in Humanoid.

704 observations, which include ball ownership information provided by the game engine, but found it to
 705 be highly inaccurate based on our visualization.

706 B.3 Ablation Study of the State Input

707 B.3.1 Vectorized States in Google Research Football

708 We perform an additional ablation study over the input of our diversity measure in GRF *3v1* scenario
 709 with SIPO-WD. We consider the following kinds of state input besides the default state input we
 710 adopted in Sec. 6:

- 711 • full observation (named *full*, 115 dims);
- 712 • default state input with random noises of the same dimension (named *random*, 36 dims).

713 The numbers of visually distinct strategies are listed in Table 8. The performance of *full* and *random*
 714 is similarly good. The result implies that the learnable discriminator can automatically filter out
 715 irrelevant states to some extent, and that SIPO-WD performs relatively robust w.r.t. different state
 input of the diversity measure.

Table 8: State input ablation. The table shows the number of distinct strategies in GRF *3v1*.

| | SIPO-WD | full | random |
|------------|-----------|-----------|-----------|
| <i>3v1</i> | 3.0 (0.0) | 3.0 (0.8) | 3.0 (0.0) |

716

717 B.3.2 RGB Images in Locomotion Tasks

718 We run SIPO-WD in the visual Humanoid task based on Isaac Gym [38]. The training protocol is
 719 similar to the state-only version (i.e., the input of policy and intrinsic rewards are both locomotion
 720 states of the Humanoid) except that we stack recent 4 RGB camera observations (84×84) as the
 721 input of intrinsic rewards in Eq. 9. We adopt the training code developed in Isaac Gym and the default
 722 PPO configuration. The backbone of the discriminator is composed of 4 convolutional layers with
 723 kernel size 3, stride 2, padding 1, and [16, 32, 32, 32] channels. Then the feature is passed to an MLP
 724 with 1 hidden layer and 256 hidden units. The activation function is leaky ReLU with slope 0.2.

725 We also compute the pairwise distance of joint torques as in the state-only version and show the
 726 result in Table 9. Visualizations are shown in Fig. 10. SIPO-WD can also learn meaningful diverse
 727 behaviors with RGB images as the state input thanks to the learnable Wasserstein discriminator. This
 728 implies that our algorithm can be naturally extended to high-dimensional states and incorporated with
 729 advances in representation learning, which may be a potential future direction.

730 B.4 How to Adjust Constraint-Related Hyperparameters

731 Three hyperparameters are essential in the implementation of the intrinsic reward r_{int} : the threshold
 732 δ , the intrinsic reward scale factor α , and the variance factor σ in $r_{\text{int}}^{\text{RBF}}$. These parameters differ under

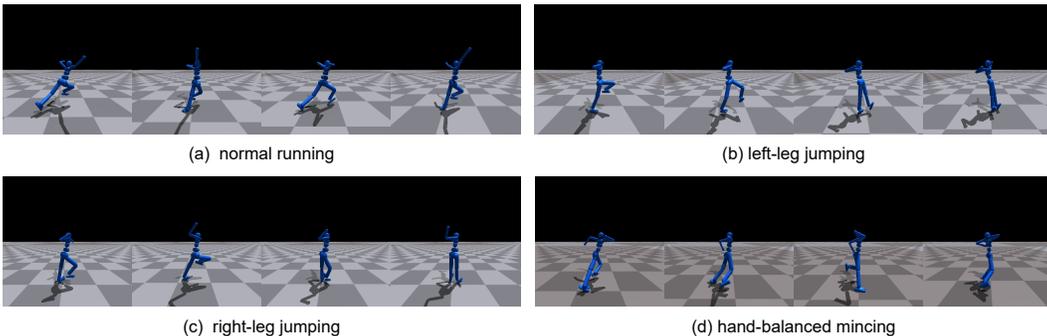


Figure 10: Results of SIPO-WD in the visual Humanoid task.

Table 9: Pairwise distance of joint torques (i.e., diversity score) in Humanoid with visual input. Results in visual experiments are averaged over 3 seeds.

| SIPO-WD (visual) | SIPO-WD | RSPO (best baseline) |
|------------------|-------------|----------------------|
| 0.62 (0.26) | 0.71 (0.23) | 0.53 (0.05) |



Figure 11: Average intrinsic reward during training π_1 .

Table 10: The values of δ and α in different environments.

| | football | | | smac | |
|-----------------------|------------|---------------|-----------|-----------------|-------------------|
| | <i>3v1</i> | <i>corner</i> | <i>CA</i> | <i>2m_vs_1z</i> | <i>2c_vs_64zg</i> |
| δ^{WD} | 0.004 | 0.01 | 0.012 | 0.02 | 0.2 |
| α^{WD} | 1 | 1 | 0.5 | 0.5 | 0.05 |
| δ^{RBF} | 0.03 | 0.01 | 0.015 | 0.002 | 0.001 |
| α^{RBF} | 0.001 | 0.001 | 0.001 | 0.001 | 0.001 |
| σ^2 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 |

733 different domains and must be adjusted individually. We find proper parameters by running two
 734 iterations without constraints and get two similar policies π_0 and π_1 . We record r_{int} during training
 735 π_1 and the trend is shown in Fig. 11. Not surprisingly, r_{int} gradually decreases as training proceeds.

736 **Threshold** We set $\delta = c_1 D_S(\pi_0, \pi_1)$. We try several different $c_1 \in \{1, 1.2, 1.4, 1.6, 1.8, 2.0\}$ and
 737 find that $c_1 = 1.2$ or 1.4 are universal proper solutions for all the experimental environments.

738 **Intrinsic Scale Factor** We need to balance the intrinsic reward r_{int} and the original reward J so
 739 that neither of the two rewards can dominate the training process. Empirically, the maximums of
 740 the two rewards should be in the same order of magnitude. i.e., $\max_{\pi} J(\pi) = \alpha \times c_2 \lambda_{\text{max}} \delta$, where
 741 $c_2 = O(1)$. When c_2 is too large, the new-trained policy π_j will oscillate near the boundary of
 742 $D(\pi_i, \pi_j) = \delta$ for some pre-trained policy p_i . Conversely, when c_2 is too small, the intrinsic reward
 743 r_{int} cannot yield diverse strategies. In experiments, we set $c_2 = 1.0$.

744 **Variance Factor** We sweep the variance factor across $\{1e-3, 5e-3, 1e-2, 2e-2, 1e-3\}$ by
 745 training π_1 and observe the trend of intrinsic rewards. We find the steepest trend and select the
 746 corresponding σ . Empirically, we find that our algorithm performs robustly well when $\sigma^2 = 0.02$.

747 The δ and α of GRF and SMAC are listed in Table 10.

748 B.5 Task Performance Evaluation

749 The evaluation win rates of the demonstrated visualization results in SMAC and GRF are shown in
 750 Table 11. Evaluated episode returns in Humanoid are shown in Table 12.

Table 11: Evaluation win rate (%) of the demonstrated visualization results in SMAC and GRF.

| | SMAC | | GRF | | |
|---------|-------------|---------------|------------|------------|---------------|
| | <i>2m1z</i> | <i>2c64zg</i> | <i>3v1</i> | <i>CA</i> | <i>corner</i> |
| π_1 | 100.0(0.0) | 98.1(2.1) | 92.3(6.2) | 48.2(10.4) | 78.2(16.2) |
| π_2 | 99.6(0.9) | 100.0(0.0) | 82.1(8.4) | 43.8(42.2) | 57.0(37.7) |
| π_3 | 100.0(0.0) | 96.9(3.3) | 90.7(1.1) | 54.7(30.6) | 55.7(20.8) |
| π_4 | 99.6(0.6) | 98.6(2.4) | 63.6(45.0) | 17.2(30.0) | 30.7(29.0) |
| π_5 | - | - | 85.4(9.1) | - | - |
| π_6 | - | - | 93.2(1.9) | - | - |
| π_7 | - | - | 64.6(32.5) | - | - |

Table 12: Episode returns in Humanoid.

| | SIPO-RBF | SIPO-WD | SIPO-WD (visual) |
|---------|---------------|---------------|------------------|
| π_1 | 4863.9(970.3) | 3909.4(533.4) | 4761.3(107.8) |
| π_2 | 3746.5(488.0) | 3784.2(481.2) | 4349.3(169.0) |
| π_3 | 3092.0(805.0) | 3770.4(674.4) | 4724.3(946.5) |
| π_4 | 2332.8(519.8) | 3589.6(387.4) | 3819.7(588.7) |

751 **B.6 Computation of Action-Based Measures in the Grid-World Example**

752 We consider the policies illustrated in Fig. 12. These policies are all optimal since these actions only
 753 include “right” and “down” and actions on non-visited states can be arbitrary. We only mark actions on
 754 on states visited by any of these 3 policies and actions on other states can be considered the same.

755 **B.6.1 Action-Distribution-Based Measures**

756 Action-distribution-based diversity measures can be defined as

$$D_{\mathcal{A}}(\pi_i, \pi_j) = \mathbb{E}_{s \sim q(s)} \left[\tilde{D}(\pi_i(\cdot | s) || \pi_j(\cdot | s)) \right], \quad (10)$$

757 where $\tilde{D}(\cdot, \cdot) : \Delta \times \Delta \rightarrow \mathbb{R}$ is a measure over action distributions and $q : \Delta(\mathcal{S})$ is a state distribution.
 758 Here, we consider q to be the joint state distribution visited by π_i and π_j .

759 **KL Divergence** KL divergence is defined by

$$D_{\text{KL}}(\pi_i(\cdot | s), \pi_j(\cdot | s)) = \int_{\mathcal{A}} \pi_i(a | s) \log \frac{\pi_i(a | s)}{\pi_j(a | s)} da.$$

760 When $\pi_j(a | s) = 0$ at any state s , KL divergence is $+\infty$. Since the trajectories of these policies
 761 have disjoint states, $D_{\mathcal{A}}^{\text{KL}}(\pi_1, \pi_2) = D_{\mathcal{A}}^{\text{KL}}(\pi_1, \pi_3) = +\infty$. Similar results can be obtained for
 762 cross-entropy.

763 **JSD $_{\gamma}$** JSD $_{\gamma}$ was defined in [34] and we consider two special cases when $\gamma = 0$ and $\gamma = 1$.

764 As illustrated by [34], JSD $_0$ measures the expected number of times two policies will “disagree”
 765 by selecting different actions. On trajectories induced by π_1 and π_2 , there are 4 + 4 states that π_1
 766 disagrees with π_2 (π_1 and π_2 are symmetric) and $D_{\mathcal{A}}^{\text{JSD}_0}(\pi_1, \pi_2) = 8/16 = 1/2$. Similarly, π_1 and
 767 π_3 only disagree at the initial state, therefore we have $D_{\mathcal{A}}^{\text{JSD}_0}(\pi_1, \pi_3) = 2/16 = 1/8$.

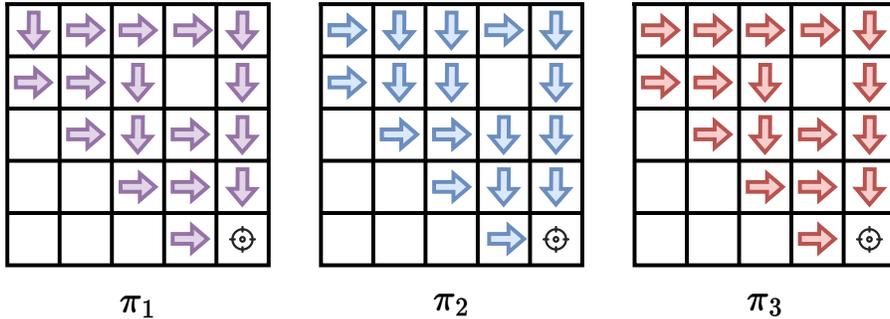


Figure 12: Policies in the grid-world example when $N_G = 5$.

768 JSD₁ is defined by

$$\begin{aligned} \text{JSD}_1(\pi_i, \pi_j) = & -\frac{1}{2} \sum_{\tau_i} P(\tau_i | \pi_i) \sum_{t=1}^T \frac{1}{T} \log \frac{\pi_i(\tau_i) + \pi_j(\tau_i)}{2\pi_i(\tau_i)} \\ & -\frac{1}{2} \sum_{\tau_j} P(\tau_j | \pi_j) \sum_{t=1}^T \frac{1}{T} \log \frac{\pi_i(\tau_j) + \pi_j(\tau_j)}{2\pi_j(\tau_j)}. \end{aligned}$$

769 Since each of the policies considered only induces a single trajectory and $\pi_i(\tau_j) = 0$ ($i \neq j$), we can
770 easily compute

$$D_{\mathcal{A}}^{\text{JSD}_1}(\pi_1, \pi_2) = D_{\mathcal{A}}^{\text{JSD}_1}(\pi_1, \pi_3) = \log 2$$

771 **Wasserstein Distance** Wasserstein distance or Earth Moving Distance (EMD) is 1 if two policies
772 disagree on a state and 0 otherwise. Therefore, it equals to $D_{\mathcal{A}}^{\text{JSD}_0}$.

773 B.6.2 Action Norm

774 We embed the action “right” as vector $[1, 0]$ since it increases the x-coordinate by 1 and the action
775 “down” as vector $[0, -1]$ since it decreases the y-coordinate by 1. This embedding can be naturally
776 extended to a continuous action space with velocity actions. Following [48], we compute the action
777 norm over a uniform distribution on states. We can see that there are 7 states where π_1 and π_2 perform
778 differently and 1 state (the initial state) where π_1 and π_3 perform differently. Therefore, we can get
779 $D(\pi_1, \pi_2) = \sqrt{7}$ and $D(\pi_1, \pi_3) = 1$.

780 B.6.3 State-Distance-Based Measures

781 **State L_2 Norm** Similar to action L_2 norm, we concatenate the coordinates instead of actions as the
782 embedding and compute the L_2 norm between embedding.

783 **Wasserstein Distance** Wasserstein distance is tractable in the grid-world example. We consider
784 7 states (except the initial and final states) in each trajectory and compute the pair-wise distance as
785 matrix C . Then we solve the following linear programming

$$\begin{aligned} \min_{\gamma} \quad & \sum_{i,j} \gamma \odot C \\ \text{s.t.} \quad & \gamma \mathbf{1} = a, \gamma^T \mathbf{1} = b \\ & \gamma_{i,j} \geq 0 \end{aligned}$$

786 where \odot means element-wise multiplication, $\mathbf{1}$ is a all-one vector, $a = [\mathbf{1}^T, \mathbf{0}^T]^T$ and $b = [\mathbf{0}^T, \mathbf{1}^T]^T$
787 is the marginal state distribution of each policy.

788 C Environment Details

789 C.1 Details of the 2D Navigation Environment

790 The navigation environment has an agent circle with size a and 4 landmark circles with size b . We
791 pre-specify a threshold c and constrain that the distance of final states reaching different landmarks
792 must be larger than c . Correspondingly, landmark circles are randomly initialized by constraining the
793 pairwise distance between centers to be larger than a threshold $c + 2(a + b)$ such that the final-state
794 constraint is valid. An episode ends if the agent touches any landmarks, i.e., the distance between the
795 center of the agent and the center of the landmark $d < a + b$, or 1000 timesteps have elapsed. The
796 observation space includes the positions of the agent and all landmarks, which is a 10-dimensional
797 vector. The action space is a 2-dimensional vector, which is the agent velocity. The time interval is
798 set to be $\Delta t = 0.1$, i.e., the next position is computed by $x_{t+1} = x_t + \Delta t \cdot v$. The reward is 1 if the
799 agent touches the landmark and 0 otherwise.

Table 13: Hyperparameters in the 2D navigation environment.

| discount | GAE λ | PPO epochs | clip parameter | entropy bonus | λ_{\max} | actor lr | critic lr | Lagrange lr | batch size |
|----------|---------------|------------|----------------|---------------|------------------|----------|-----------|-------------|------------|
| 0.997 | 0.95 | 10 | 0.2 | 0 | 10 | 3e-4 | 1e-3 | 0.5 | 4000 |

800 C.2 Details of Environments

801 We provide training configurations and environment introductions below and refer readers to our
802 project website in App. A for visualizations of these environments.

803 **Humanoid** We use the Humanoid environment in IsaacGym [38] with default observation and
804 action spaces. The input of intrinsic rewards or diversity measure is the observation without all torque
805 states.

806 **SMAC** We adopt the SMAC environment in the MAPPO codebase¹ with the same configuration
807 as Yu et al. [64]. The input of intrinsic rewards or diversity measure is the state of all allies, including
808 positions, health, etc.

809 On the “easy” map *2m_vs_1z*, two marines must be controlled to defeat a Zealot. The marines can
810 attack from a distance, while the Zealot’s attacks are limited to close range. A successful strategy
811 involves alternating the marines’ attacks to distract the Zealot. On the “hard” map *2c_vs_64zg*, two
812 colossi must be controlled by the agents to fight against 64 zergs. The colossi have a wider attack
813 range and can move over cliffs. Strategies on this map may include hit-and-run tactics, waiting in
814 corners, or dividing and conquering enemies. The level of difficulty is determined by the learning
815 performance of existing MARL algorithms. Harder maps require more exploration and training steps.

816 **GRF** We adopt the “simple115v2” representation as observation with both “scoring” and “check-
817 point” reward. The reward is shared across all agents. The input of intrinsic rewards or diversity
818 measure is the position and velocity of all attackers and the ball. All policies are trained to control
819 the left team to score against build-in bots.

820 *academy_3_vs_1_with_keeper*: In this scenario, a team of three players (left) try to score a goal
821 against a single defender and a goalkeeper. The left team starts with the ball and has to dribble past
822 the defender and the goalkeeper to score a goal.

823 *academy_counterattack_easy*: In this scenario, the left team starts with the ball in the front-yard and
824 try to score a goal against several defenders. All eleven players in the left players can be controlled.

825 *academy_corner*: In this scenario, the left team tries to score a goal from a corner kick. The right
826 team defends the goal and tries to prevent the left team from scoring. All eleven players in the left
827 players can be controlled.

828 D Implementation Details

829 D.1 2D Navigation

830 We apply PPO with Lagrange multipliers to optimize the policy and hyperparameters are summarized
831 in Table 13. $D(\pi_i, \pi_j)$ is simply taken as the L_2 distance of the final state reached by π_i and π_j , i.e.,
832 $D(\pi_i, \pi_j) = \|s_H^{\pi_i} - s_H^{\pi_j}\|^2$. The applied algorithm is the same as SIPO (see Appendix G) except that
833 the intrinsic reward is only computed at the last timestep.

834 D.2 SIPO

835 In the i -th iteration ($1 \leq i \leq M$), we learn an actor and a critic with i separate value heads to
836 accurately predict different return terms, including $i - 1$ intrinsic returns for the diversity constraints
837 and the environment reward. We include all practical tricks mentioned in [64] because we find
838 them all critical to algorithm performance. We use separate actor and critic networks, both with
839 hidden size 64 and a GRU layer with hidden size 64. The common hyperparameters for SIPO,

¹<https://github.com/marlbenchmark/on-policy>

Table 14: Common hyperparameters for SIPO, baselines, and ablations.

| discount | GAE λ | actor lr | critic lr | clip parameter | entropy bonus | GRF batch size | SMAC batch size |
|----------|---------------|----------|-----------|----------------|---------------|----------------|-----------------|
| 0.99 | 0.95 | 5e-4 | 1e-3 | 0.2 | 0.01 | 9600 | 3200 |

Table 15: SIPO hyperparameters across all environments.

| λ_{\max} | Discriminator lr | Lagrangian lr |
|------------------|------------------|---------------|
| 10 | 4.0e-4 | 0.1 |

840 baselines, and ablations are listed in Table 14. Other environment-specific parameters, such as PPO
 841 epochs and mini-batch size, are all the same as [64]. Besides, Table 10 and Table 15 lists some extra
 842 hyperparameters for SIPO.

843 D.3 Baselines

844 We re-implement all baselines with PPO based on the MAPPO [64] project. All algorithms run for
 845 the same number of environment frames. Specific hyperparameters for baselines can be found in
 846 Appendix D.3.

847 **SMERL** SMERL trains a latent-conditioned policy that can robustly adapt to new scenarios. It
 848 promotes diversity by maximizing the mutual information between states and the latent variable.
 849 We implement SMERL with PPO, where the actor and the critic take as the input the concatenation
 850 of observation and a one-hot latent variable. The discriminator is a 2-layer feed-forward network
 851 with 64 hidden units. The learning rate of the discriminator is the same as the learning rate of
 852 the critic network. The input of the discriminator is the same as the input we use for SIPO-WD.
 853 The critic has 2 value heads for an accurate estimation of intrinsic return. Since SMERL trains a
 854 single latent-conditioned policy, we train SMERL for $M \times$ more environment steps, such that total
 855 environment frames are the same. The scaling factor of intrinsic rewards is 0.1 and the threshold
 856 for diversification is [0.81, 0.45, 0.72] ($0.9 \times [0.9, 0.5, 0.8]$) for “3v1”, “counterattack”, and “corner”
 857 respectively.

858 **DvD** DvD simultaneously trains a population of policies to maximize the determinant of a kernel
 859 matrix based on action difference. We concatenate the one-hot actions along a trajectory as the
 860 behavioral embedding. The square of the variance factor, i.e., σ^2 in the RBF kernel, is set to be the
 861 length of behavioral embedding. We also use the same Bayesian bandits as proposed in the original
 862 paper. Training DvD in “counterattack” and “corner” exceeds the GPU memory and we exclude the
 863 results in the main body.

864 **DIPG** DIPG iteratively maximizes the maximum mean discrepancy (MMD) distance between the
 865 state distribution of the current policy and previously discovered policies. For DIPG, we follow the
 866 opensource implementation². We set the same variance factor in the RBF kernel as SIPO-RBF and
 867 apply the same state as the input of the RBF kernel. We sweep the coefficient of MMD loss among
 868 $\{0.1, 0.5, 0.9\}$ and find 0.1 the most appropriate (larger value will cause training instability). We use
 869 the same method to save archived trajectories as SIPO and the input of the RBF kernel is the same as
 870 the input we use for SIPO-RBF. To improve training efficiency, we only back-propagate the MMD
 871 loss at the first PPO epoch.

872 **RSPO** RSPO iteratively discovers diverse policies by optimizing extrinsic rewards on novel tra-
 873 jectories while optimizing diversity on other trajectories. The diversity measure is defined as the
 874 action-cross entropy along the trajectory. For RSPO, we follow the opensource implementation³
 875 and use the same hyperparameters on the SMAC *2c_vs_64zg* map in the original paper for GRF
 876 experiments.

²<https://github.com/dtak/DIPG-public>

³<https://github.com/footoredo/rsपो-iclr-2022>

877 **TrajDi** TrajDi was originally designed for cooperate multi-agent domains to facillitate zero-shot
878 coordination. It defines a generalized Jensen-Shanon divergence objective between policy action distri-
879 butions. Then this objective and rewards are simultaneously optimized via population-based training.
880 We tried TrajDi in SMAC and GRF. We sweep the action discount factor among $\{0.1, 0.5, 0.9\}$ and
881 the coefficient of TrajDi loss among $\{0.1, 0.01, 0.001\}$. However, TrajDi fails to converge in the “3v1”
882 scenario and exceeds the GPU memory in the “counterattack” and “corner” scenarios. Therefore, we
883 exclude the performance of TrajDi in the main body.

884 D.4 Ablation Study Details

885 For the three ablation studies: fix-L, CE, and filter, we list the specific hyperparameters here:

- 886 • fix-L: we set the Lagrange multiplier to be 0.2;
- 887 • CE: the threshold is 3.800 and the intrinsic reward scale factor is 1/1000 of that in the WD
888 setting;
- 889 • filter: all the hyperparameters in the setting is the same as those in the WD setting.

890 E Proofs

891 E.1 Proof of theorem 4.1

892 **Theorem 4.1.** Assume D is a distance metric. Denote the optimal value of Problem 3 as T_1 . Let
893 $T_2 = \sum_{i=1}^M J(\tilde{\pi}_i)$ where

$$\begin{aligned} \tilde{\pi}_i &= \arg \max_{\pi_i} J(\pi_i) \\ \text{s.t. } D(\pi_i, \tilde{\pi}_j) &\geq \delta/2, \quad \forall 1 \leq j < i \end{aligned} \quad (3)$$

894 for $i = 1, \dots, M$, then $T_2 \geq T_1$.

895 *Proof.* Suppose the optimal solution of Problem 3 is $\pi_1, \pi_2, \dots, \pi_M$ satisfying $J(\pi_1) \geq J(\pi_2) \geq$
896 $\dots \geq J(\pi_M)$ and the optimal solution of Problem 6 is $\tilde{\pi}_1, \tilde{\pi}_2, \dots, \tilde{\pi}_M$ satisfying $J(\tilde{\pi}_1) \geq J(\tilde{\pi}_2) \geq$
897 $\dots \geq J(\tilde{\pi}_M)$.

Assume the contrary that Thm. 4.1 is not true, which means $\sum_{i=1}^M J(\pi_i) = T_1 > T_2 = \sum_{i=1}^M J(\tilde{\pi}_i)$.
Then we choose the smallest number $N \leq M$ that satisfies

$$\sum_{i=1}^N J(\pi_i) > \sum_{i=1}^N J(\tilde{\pi}_i).$$

898 By $T_1 > T_2$ we know that N exists. In addition, because Problem 6 solves unconstrained RL in the
899 first iteration, we know that $\tilde{\pi}_1 = \arg \max_{\pi} J(\pi)$ and then $J(\pi_1) \leq J(\tilde{\pi}_1)$. Therefore, $N \geq 2$.

900 Suppose $J(\pi_N) \leq J(\tilde{\pi}_N)$. Then we have

$$\sum_{i=1}^{N-1} J(\pi_i) > \sum_{i=1}^{N-1} J(\tilde{\pi}_i).$$

901 Contradicting the fact that N is the smallest number satisfies that equation.

Hence, we know that $J(\pi_N) > J(\tilde{\pi}_N)$. Then

$$J(\pi_1) \geq J(\pi_2) \geq \dots \geq J(\pi_N) > J(\tilde{\pi}_N).$$

902 Consider the optimization problem of $\tilde{\pi}_N$:

$$\begin{aligned} \tilde{\pi}_N &= \arg \max_{\pi} J(\pi) \\ \text{s.t. } D(\pi, \tilde{\pi}_j) &\geq \delta/2, \quad \forall 1 \leq j < N. \end{aligned}$$

903 This optimization does not find $\{\pi_1, \dots, \pi_N\}$ but find $\tilde{\pi}_N$, which means that for each π_i , $1 \leq i \leq N$,
 904 there exists $1 \leq j_i < N$ such that $D(\pi_i, \tilde{\pi}_{j_i}) < \delta/2$. Otherwise, we will get the solution of the above
 905 problem as π_i instead of $\tilde{\pi}_N$.

By the Pigeonhole Principle, we know that there exist two indexes $i_1 \in [N]$ and $i_2 \in [N]$ ($i_1 \neq i_2$)
 such that $j_{i_1} = j_{i_2} = \hat{j}$. Then we have

$$D(\pi_{i_1}, \pi_{i_2}) \leq D(\pi_{i_1}, \tilde{\pi}_{\hat{j}}) + D(\pi_{i_2}, \tilde{\pi}_{\hat{j}}) < \delta/2 + \delta/2 = \delta,$$

906 where the inequality follows by the triangle inequality of the distance function.

907 It contradict with the fact that $D(\pi_{i_1}, \pi_{i_2}) \geq \delta$ in Problem 3.

908 Therefore, we prove the theorem $\sum_{i=1}^M J(\pi_i) = T_1 \leq T_2 = \sum_{i=1}^M J(\tilde{\pi}_i)$. \square

909 E.2 Proof of Theorem 5.1

910 In this section, we consider the i -th iteration of SIPO illustrated in Eq. (4). For the sake of simplicity,
 911 we use $a \leq \boldsymbol{\lambda} \leq b$ for vector $\boldsymbol{\lambda}$ to denote each component of $\boldsymbol{\lambda}$ satisfies $a \leq \lambda_i \leq b$, where $a, b \in \mathbb{R}$.
 912 We use π to denote the policy we are optimizing, and π_j ($1 \leq j < i$) to denote a previously obtained
 913 policy. We denote the Lagrange function as $L(\pi, \boldsymbol{\lambda}) = -J(\pi) - \sum_{j=1}^{i-1} \lambda_j (D(\pi, \pi_j) - \delta)$.

914 To prove Theorem 5.1, we consider the following two optimization problems:

$$(\pi_i, \boldsymbol{\lambda}^*) = \arg \min_{\pi} \max_{\boldsymbol{\lambda} \geq 0} L(\pi, \boldsymbol{\lambda}) \quad (11)$$

915 and

$$(\tilde{\pi}_i, \tilde{\boldsymbol{\lambda}}^*) = \arg \min_{\pi} \max_{0 \leq \lambda \leq \Lambda} L(\pi, \boldsymbol{\lambda}), \quad (12)$$

916 where $\Lambda = \frac{1}{\epsilon_0}$ and $\epsilon_0 > 0$ is sufficiently small.

917 We make the following assumptions to prove this theorem:

918 **Assumption E.1.** $0 \leq J(\cdot) \leq 1$.

919 **Assumption E.2.** $\forall \boldsymbol{\lambda} \geq 0$, $L(\cdot, \boldsymbol{\lambda})$ is l -smooth and ζ -Lipschitz.

920 We may notice that, solving optimization problem (11) is hard because its domain is unbounded.
 921 Therefore, we make some approximation and consider bounded optimization problem (12). First we
 922 prove the following lemma about the value function J :

923 **Lemma E.3.** $J(\pi_i) \leq J(\tilde{\pi}_i)$.

924 *Proof.* As the domain of $\boldsymbol{\lambda}$ in Eq. 12 is smaller than Eq. (11), we have $L(\pi_i, \boldsymbol{\lambda}) \geq L(\tilde{\pi}_i, \tilde{\boldsymbol{\lambda}})$.

925 By the fundamental property of Lagrange duality, we know that L achieves its optimal value when
 926 $\boldsymbol{\lambda} = 0$ and the optimal value is $-J(\pi_i)$.

927 By the optimality of $(\tilde{\pi}_i, \tilde{\boldsymbol{\lambda}}^*)$, we know that

$$-\sum_{j=1}^{i-1} \tilde{\lambda}_j^* (D(\tilde{\pi}_i, \pi_j) - \delta) \geq 0. \quad (13)$$

Then we have

$$-J(\pi_i) = L(\pi_i, \boldsymbol{\lambda}^*) \geq \tilde{L}(\tilde{\pi}_i, \tilde{\boldsymbol{\lambda}}^*) = -J(\tilde{\pi}_i) - \sum_{j=1}^{i-1} \tilde{\lambda}_j^* (D(\tilde{\pi}_i, \pi_j) - \delta) \geq -J(\tilde{\pi}_i).$$

928 \square

929 Then we prove the distance between optimal policy $\tilde{\pi}_i$ in problem (12) and optimal policy π_i in
 930 problem (11) is very small:

931 **Lemma E.4.** Under Assumption E.1, $D(\tilde{\pi}_i, \pi_j) \geq \delta - \epsilon_0$, $\forall 1 \leq j < i$.

932 *Proof.* We prove by contradiction.

933 Suppose there exists $1 \leq j_0 < i$, $D(\tilde{\pi}_i, \pi_{j_0}) < \delta - \epsilon_0$. Then we choose $\hat{\lambda}$ such that

$$\hat{\lambda}_j = \begin{cases} \Lambda & j = j_0, \\ 0 & 1 \leq j < i, j \neq j_0. \end{cases}$$

934 By the Assumption E.1, Eq. (13), and $\Lambda = \frac{1}{\epsilon_0}$, we have

$$0 \geq -J(\pi_i) = L(\pi_i, \boldsymbol{\lambda}^*) \geq L(\tilde{\pi}_i, \tilde{\boldsymbol{\lambda}}^*) \geq L(\tilde{\pi}_i, \hat{\boldsymbol{\lambda}}) \geq -1 - \Lambda(D(\tilde{\pi}_i, \pi_{j_0}) - \delta) > 0.$$

That is a contradiction. So we have proved that

$$D(\tilde{\pi}_i, \pi_j) \geq \delta - \epsilon_0, \quad \forall 1 \leq j < i.$$

935

□

936 From the deduction above, we get the following approximation lemma:

937 **Lemma E.5.** Denote the optimal solution of Eq. 11 and Eq. 12 as (π_i, λ) and $(\tilde{\pi}_i, \tilde{\lambda})$ respectively.
938 Then we have the following approximation about the optimal value and distance:

$$\begin{aligned} J(\pi_i) &\leq J(\tilde{\pi}_i) \\ D(\tilde{\pi}_i, \pi_j) &\geq \delta - \epsilon_0, \quad \forall 1 \leq j < i \end{aligned}$$

939 *Proof.* This lemma follows directly by Lemma E.3 and Lemma E.4. □

940 Therefore, it is reasonable to consider the constrained optimization problem (12) instead of primal
941 problem (11) because we have proved that the optimal value doesn't get smaller and the distance of
942 policy is ϵ_0 -approximation of the primal problem. Finally we use the conclusion in the paper [27] to
943 analysis the convergence of problem (12):

Lemma E.6. ([27], Theorem 4.8) Under Assumption E.2, solving Eq. (12) via two-timescale GDA with learning rate $\eta_\pi = \Theta(\epsilon^4/l^3\zeta^2\Lambda^2)$ and $\eta_\lambda = \Theta(1/l)$ requires

$$\mathcal{O}\left(\frac{l^3\zeta^2\Lambda^2C_1}{\epsilon^6} + \frac{l^3\Lambda^2C_2}{\epsilon^4}\right)$$

944 iterations to converge to an ϵ -stationary point π_i^* , where C_1 and C_2 are the constants that depend on
945 the distance between the initial point and the optimal point.

946 **Theorem 5.1.** Under assumptions E.1 and E.2 and learning rate with learning rate $\eta_\pi =$
947 $\Theta(\epsilon^4/l^3\zeta^2\Lambda^2)$ and $\eta_\lambda = \Theta(1/l)$, SIPO converges to an ϵ -stationary point with convergence rate
948 $\mathcal{O}\left(\frac{l^3\zeta^2\Lambda^2C_1}{\epsilon^6} + \frac{l^3\Lambda^2C_2}{\epsilon^4}\right)$.

949 *Proof.* We consider the following constraint nonconvex-concave optimization:

$$\min_{\pi} \max_{0 \leq \boldsymbol{\lambda} \leq \Lambda} L(\pi, \boldsymbol{\lambda}). \quad (14)$$

950 Following Lemma E.6, we know that the Two-Timescale GDA algorithm converges to an ϵ -stationary
951 point π_i^* .

952 From the above deduction, the Two-Timescale GDA algorithm requires $\mathcal{O}\left(\frac{l^3\zeta^2\Lambda^2C_1}{\epsilon^6} + \frac{l^3\Lambda^2C_2}{\epsilon^4}\right)$
953 iterations with learning rate $\eta_\pi = \Theta(\epsilon^4/l^3\zeta^2\Lambda^2)$ and $\eta_\lambda = \Theta(1/l)$ to converge to an ϵ -stationary
954 point with convergence rate.

955

□

956 **F Discussion**

957 **F.1 The Failure Case of State-Distance-Based Diversity Measures**

958 A failure case of state-distance-based diversity measures may be when the state space includes
959 many *irrelevant features*. These features cannot reflect behavioral differences. If we run SIPO
960 in such an environment, the learned strategies may be only diverse w.r.t these features and have
961 little visual distinction. Like the famous noisy TV problem [4], the issue of irrelevant features is
962 intrinsically challenging for general RL applications, which cannot be resolved by using action-based
963 or state-occupancy-based diversity measures either.

964 Thanks to the advantages we discussed in the paper, we generally find that state-distance-based
965 measures can be preferred in challenging RL problems. Meanwhile, since the state dimension can be
966 much higher than actions, it is possible that RL optimization over states may be accordingly more
967 difficult than actions. In practice, we can design a feature selector for those most relevant features
968 for visual diversity and run diversity learning over the filtered features. In SMAC and GRF, we
969 utilize the agent features (excluding enemies) as the input of diversity constraint without further
970 modifications, as discussed in Appendix D. We remark that even after filtering, the agent features
971 remain high-dimensional while our algorithm still works well. Note that using a feature selector is
972 a common practice in many existing domains, such as novelty search [11], exploration [29], and
973 curriculum learning [6]. There are also works studying how to extract useful low-dimensional features
974 from observations [63, 16], which are orthogonal to our focus.

975 **F.2 The Distance Metric**

976 In Sec. 5, we adopt the two most popular implementations in the machine learning literature, i.e.,
977 RBF kernel and Wasserstein distance, while it is totally fine to adopt alternative implementations.
978 For example, we can learn state representations (e.g. auto-encoder, Laplacian, or successor feature)
979 and utilize pair-wise distance or norms as a diversity measure. Similar topics have been extensively
980 discussed in the exploration literature [63, 37]. We leave them as our future directions.

981 **G Pseudocode of SIPO**

982 The pseudocode of SIPO is shown in Algorithm 1.

Algorithm 1 SIPO (red for SIPO-RBF and blue for SIPO-WD)

Input: Number of Iterations M , Number of Training Steps within Each Iteration T .
Hyperparameter: Learning Rate η_π , Diversity Threshold δ , Intrinsic Scale Factor α , Lagrange Multiplier Upperbound λ_{\max} , Lagrange Learning rate η_λ , Wasserstein Critic Learning Rate η_W , RBF Kernel Variance σ .

- 1: Archived trajectories $X \leftarrow \emptyset$ // to store states visited by previous policies
- 2: **for** iteration $i = 1, \dots, M$ **do**
- 3: Initialize policy π_{θ_i} // initialization
- 4: Initialize Wasserstein critic f_{ϕ_i}
- 5: **for** archive index $j = 1, \dots, i - 1$ **do**
- 6: Lagrange multiplier $\lambda_j \leftarrow 0$
- 7: **end for**
- 8: **for** Training step $t = 1, \dots, T$ **do**
- 9: Collect trajectory $\tau = \{(s_h, \mathbf{a}_h, r(s_h, \mathbf{a}_h))\}_{h=1}^H$
- 10: **for** archive index $j = 1, \dots, i - 1$ **do**
- 11: $R_{\text{int}}^j \leftarrow 0$
- 12: **end for**
- 13: **for** timestep $h = 1, \dots, H$ **do** // compute intrinsic reward
- 14: $r_{\text{int},h} \leftarrow 0$
- 15: **for** archive trajectory $\chi_j \in X$ **do**
- 16: $r_{\text{int},h}^j \leftarrow -\frac{1}{H|\chi_j|} \sum_{s' \in \chi_j} \exp\left(-\frac{\|s_h - s'\|^2}{2\sigma^2}\right)$
- 17: $r_{\text{int},h}^j \leftarrow \frac{1}{H} \left[f_{\phi_j}(s_h) - \frac{1}{|\chi_j|} \sum_{s' \in \chi_j} f_{\phi_j}(s') \right]$
- 18: $r_{\text{int},h} \leftarrow r_{\text{int},h} + \lambda_j \cdot r_{\text{int},h}^j$
- 19: $R_{\text{int}}^j \leftarrow R_{\text{int}}^j + r_{\text{int},h}^j$
- 20: **end for**
- 21: $r_h \leftarrow r(s_h, \mathbf{a}_h) + \alpha \cdot r_{\text{int},h}$
- 22: **end for**
- 23: **for** archive index $j = 1, \dots, i - 1$ **do**
- 24: $\lambda_j \leftarrow \text{clip}\left(\lambda_j + \eta_\lambda \left(-R_{\text{int}}^j + \delta\right), 0, \lambda_{\max}\right)$ // gradient ascent on λ_j
- 25: $\phi_j \leftarrow \phi_j + \eta_W \frac{1}{H} \sum_{h=1}^H \nabla_{\phi_j} \left(f_{\phi_j}(s_h) - \frac{1}{|\chi_j|} \sum_{s' \in \chi_j} f_{\phi_j}(s') \right)$
- 26: $\phi_j \leftarrow \text{clip}(\phi_j, -0.01, 0.01)$
- 27: **end for**
- 28: Update π_{θ_i} with $\{(s_h, \mathbf{a}_h, r_h)\}$ by PPO algorithm // policy gradient on θ_i
- 29: **end for**
- 30: Collect many trajectories χ_i // collect trajectories to approximate $d_{\pi_{\theta_i}}$
- 31: $X \leftarrow X \cup \{\chi_i\}$ // for the use of following iterations
- 32: **end for**
