

Open-Ended Diverse Solution Discovery with Regulated Behavior Patterns for Cross-Domain Adaptation

Kang Xu, Yan Ma, Bingsheng Wei, Wei Li *

Academy for Engineering and Technology, Fudan University, Shanghai, China
kangxu21@m.fudan.edu.cn, {20210860024, weibingsheng, fd_liwei}@fudan.edu.cn

Abstract

While Reinforcement Learning can achieve impressive results for complex tasks, the learned policies are generally prone to fail in downstream tasks with even minor model mismatch or unexpected perturbations. Recent works have demonstrated that a policy population with diverse behavior characteristics can generalize to downstream environments with various discrepancies. However, such policies might result in catastrophic damage during the deployment in practical scenarios like real-world systems due to the unrestricted behaviors of trained policies. Furthermore, training diverse policies without regulation of the behavior can result in inadequate feasible policies for extrapolating to a wide range of test conditions with dynamics shifts. In this work, we aim to train diverse policies under the regularization of the behavior patterns. We motivate our paradigm by observing the inverse dynamics in the environment with partial state information and propose *Diversity in Regulation* (DiR) training diverse policies with regulated behaviors to discover desired patterns that benefit the generalization. Considerable empirical results on various variations of different environments indicate that our method attains improvements over other diversity-driven counterparts.

Introduction

Deep Reinforcement Learning has exhibited wide success in solving complex tasks, including vision-based video games (Mnih et al. 2015; Jaderberg et al. 2019), quadruped locomotion (Hwangbo et al. 2019; Lee et al. 2020a), and robotic manipulation (Andrychowicz et al. 2020). However, the policies trained in the source environments are prone to fail in environmental variations. For example, dynamics change, such as the damaged component of a robot or encountering a new terrain, might lead to a failure due to poor generalization of the agents. Thus one significant challenge for real-world deployment of RL is the generalization across various conditions.

One natural approach is to train a policy under a range of dynamics in simulation (Tobin et al. 2017; Peng et al. 2018; Rajeswaran et al. 2016), and assume that the trained policy can generalize to the specific target environment.

These methods require expert knowledge to manually specify the distribution of training environments in a trial-and-error manner to involve the properties of the target environment. In addition, the trained policy may appear to be over-conservative due to the uncertainty in the training environments (Yu, Liu, and Turk 2018; Xie et al. 2021). Another category of methods trains the policy to implicitly identify the dynamics of the target environment based on samples collected in the target environment during training and then encourages the policy to act optimally according to the identified dynamics (Muratore et al. 2021; Du et al. 2021; Evans, Thankaraj, and Pinto 2022). However, rollouts in the target environment, like the real-world system during training, might result in catastrophic damage due to the premature behavior of the policy.

Recent works have demonstrated that diversity-driven policies can extrapolate to new environments through the few-shot adaptation (Eysenbach et al. 2018; Kumar et al. 2020; Osa, Tangkaratt, and Sugiyama 2021; Parker-Holder et al. 2020; Zhou et al. 2022). While the policy population with different behavior characteristics can generalize to different environment variations, the learned policies may result in potential safety problems in practical scenarios like real-world systems, as the behaviors of the diverse policies are unpredictable. Especially, the degree of diversity that is necessary for the generalization may be limited. For instance, when we aim to obtain multiple policies for a quadruped robot that can generalize to various terrains, what we desire might be policies with different locomotion patterns rather than those able to roll on the ground. However, these works train the policies without regularizing the behavior, which might result in inadequate feasible policies.

In this work, we take the first step towards diverse policies with regulated behaviors for generalization. To encourage sufficient feasible solutions for adaptation in a wide range of downstream scenarios, we propose a novel diversity objective based on the divergence of inverse dynamics models $\mathcal{T}^\pi(a|f(s), f(s'))$ under partial state information. The partial state information is removed by utilizing a customizable state filtration function $f(s)$. Intuitively, the actions impacting the remaining state information would be discouraged from getting diversified, thus regulating the behaviors of trained policies. Additionally, we introduce the

*Corresponding Author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

open-ended training manner to achieve continuous solution discovery, which avoids the drawback of prior work training with a fixed number of policies (Kumar et al. 2020; Parker-Holder et al. 2020).

The main contribution of our work is the proposal of a diversity-driven algorithm, *Diversity in Regulation* (DiR), which trains multiple policies with regulated behavior patterns for efficient generalization by diversifying the action distributions in a customizable way. Our analysis demonstrates that the discovered policies show more regulated behaviors against prior diversity-driven approaches, which benefits generalization across a wide range of test conditions. Empirically, we observe that DiR substantially outperforms prior methods under various environment discrepancies.

Preliminaries

Notation. To model the sequential decision problem, we consider the standard Markov Decision Process (MDP) (Sutton and Barto 2018) as $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \mu, \gamma)$, where \mathcal{S} and \mathcal{A} are the state space and action space respectively; $\mathcal{P}(s'|s, a) : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ specifies the dynamics of the environment and defines the transition probability of reaching s' at the next step given current state s and the executed action a ; $r(s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$ denotes the reward function; $\mu(s) : \mathcal{S} \rightarrow [0, 1]$ denotes the distribution of initial states and $\gamma \in (0, 1)$ is the discount factor. Considering a policy $\pi(a|s) : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ which outputs the probability of choosing action a given the state s , the probability density function of any trajectory $\tau = \{s_0, a_0, s_1, a_1, \dots, s_T\}$ can be formulated as $P(\tau) = \mu(s_0) \prod_{t=0}^{T-1} \pi(a_t|s_t) \mathcal{P}(s_{t+1}|s_t, a_t)$.

Inverse dynamics. Here we denote the inverse dynamics of the MDP as $\mathcal{T}(a|s, s') : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ which defines the probability of action a given the state pair (s, s') in the consecutive steps. Since there might be various actions under different policies given the state pair (s, s') , the inverse dynamics under some specific policy π can be formulated as:

$$\mathcal{T}^\pi(a|s, s') = \frac{\mathcal{P}(s'|s, a)\pi(a|s)}{\int_{\mathcal{A}} \mathcal{P}(s'|s, \hat{a})\pi(\hat{a}|s)d\hat{a}}. \quad (1)$$

State filtration function. In this work, we assume a customizable function $f(s) : \mathcal{S} \rightarrow \tilde{\mathcal{S}}$ that removes partial state information from the state s . For instance, $f(s)$ can be defined to remove the x-axis coordinate in a navigation task whose full state information includes 2D coordinates. The resulting partial state space $\tilde{\mathcal{S}}$ can also be considered as the state space from a partially observable MDP (POMDP) (Bellman 1957).

Mutual-Information in RL. Mutual information can be generally expressed as

$$\begin{aligned} I(X; Y) &= \iint p(x, y) \log \frac{p(x, y)}{p(x)p(y)} dy dx \\ &= H(X) - H(X|Y) = H(Y) - H(Y|X), \quad (2) \end{aligned}$$




Figure 1: Left: Inference variance of inverse dynamics models trained under different state settings in Walker2D increases with the increase of missing state information. Right: Toy example of a 2-DOF robotic arm to interpret the inference variance of IDMs under different state settings. Given the partial state information pair (\bar{s}, \bar{s}') , there will be more possible actions compared with the full state information setting.

which defines the mutual dependence between two random variables. Mutual information has been introduced to find the best representation subject to a constraint on the complexity (Tishby, Pereira, and Bialek 2000; Alemi et al. 2016). In the context of RL, maximizing the mutual information $I(S; Z) = H(Z) - H(Z|S)$ between visited states S and latent variable Z has been proposed to discover diverse policies (Eysenbach et al. 2018; Kumar et al. 2020).

Diverse high-performing policies. Prior approaches that train multiple high-performing policies $\{\pi_{\theta_k}\}_{k=1}^M$ with diverse behaviors¹ can be naturally formulated as:

$$\begin{aligned} \max_{\pi_k} E_{\pi_k} \left[\sum_{t=1}^T \gamma^{t-1} r_t \right], \quad \forall 1 \leq k \leq M \quad (3) \\ \text{s.t. } D(\pi_i, \pi_j) \geq \delta, \quad \forall 1 \leq i, j \leq M, i \neq j, \end{aligned}$$

where $D(\cdot, \cdot)$ measures how different the two policies are, and δ is the diversity threshold. One natural choice of the distance measurement function is the KL divergence $D_{KL}[\pi_i(\cdot|s) \parallel \pi_j(\cdot|s)]$ which is widely adopted in prior works (Schulman et al. 2015, 2017; Hong et al. 2018).

Few-shot adaptation. In novel environments, we rollout each policy from the trained population for several episodes and deploy the best-performing one, which resembles prior diversity-driven approaches (Kumar et al. 2020; Osa, Tangkaratt, and Sugiyama 2021; Derek and Isola 2021).

Motivation Example: Inference Variance of Inverse Dynamics Models

We motivate our method with an empirical observation about the inference variance of the inverse dynamics under different state information settings. Here we train four independent inverse dynamics models (IDMs) $\{\mathcal{T}_{\phi_i}(a|s, s')\}_{i=1}^4$ simultaneously using different state settings in Walker from Mujoco (Todorov, Erez, and Tassa 2012). The training data are collected by an agent trained

¹The θ will be omitted for simplicity.

through Soft Actor-Critic (SAC) (Haarnoja et al. 2018). We utilized four different state settings for the inverse dynamics models with varying degrees of information missing. We train the inverse dynamics models by maximizing the log-likelihood:

$$L(\phi_i) = E_{(s,a,s') \sim \mathcal{D}} [\log \mathcal{T}_{\phi_i}(a|f_i(s), f_i(s'))], \quad 1 \leq i \leq 4, \quad (4)$$

where the $f_i(s) : \mathcal{S} \rightarrow \bar{\mathcal{S}}_i$ is the state filtration function corresponding to the model i that removes some specific state information. After training, we evaluate the inference variance $\text{Var}[\mathcal{T}_{\phi_i}(\cdot|s, s')]$ of the IDMs given the same batch of data $\mathcal{D}_{test} = \{(s, s')\}$, and the results are shown in Figure 1. The results indicate that the inference variance of the inverse dynamics model increases with the increase of missing state information. The details of the experiment refer to Appendix A.

Herein we interpret the observation with a toy example of a 2-DOF robotic arm, as shown in Figure 1. When we employ the joint angle θ_1 as a partial state space $\bar{\mathcal{S}}$, there might be multiple possible actions given any state pair (\bar{s}, \bar{s}') . In contrast, there might be only one possible action given the state pair (s, s') under the full state information setting. Motivated by the observation, we propose to maximize the divergence of IDMs with the partial state setting under different policies to encourage them to produce distinct action distributions given any $(f(s), f(s'))$ pair. Furthermore, we introduce customizable state filtration functions to specify the desired patterns. Semantically, the removing state information (e.g., θ_2 above) can be regarded as the state information of body parts that is unnecessary to be diversified (e.g., roll angles of a quadruped robot). In contrast, we can remove the state information like leg motions to obtain policies with diverse locomotion patterns. Thus, we introduce the state filtration function to focus on discovering desired behaviors.

We denote the divergence between inverse dynamics under the partial state information setting as $P^{\pi_j}(a|f(s), f(s')) : \bar{\mathcal{S}} \times \bar{\mathcal{S}} \times \mathcal{A} \rightarrow [0, 1]$, and the overall objective can be formulated as the following constrained optimization problem:

$$\max_{\pi_k} E_{\pi_k} \left[\sum_{t=1}^T \gamma^{t-1} r_t \right], \quad \forall 1 \leq k \leq M \quad (5)$$

$$\text{s.t. } E[\text{D}_{KL}[P^{\pi_i}(\cdot|f(s), f(s')) \| P^{\pi_j}(\cdot|f(s), f(s'))]] \geq \delta, \\ \forall 1 \leq i, j \leq M, i \neq j,$$

where $f(s) : \mathcal{S} \rightarrow \bar{\mathcal{S}}$ removes specific partial state information from s , and can be designed to control the diverse patterns we aim to discover.

Diversity in Regulation

This section presents our approach to resolving the objective in detail. We propose an open-ended training manner for diverse solution discovery, transform the objective into a trivial form, and finally analyze the connection between our method and prior diversity-driven approaches through mutual information. The overview of DiR is shown in Figure 2.

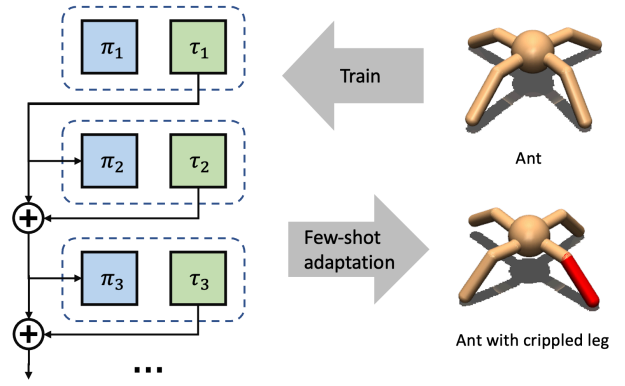


Figure 2: The semantic overview of DiR. π_i and τ_i represent the policy and inverse dynamics model, respectively. We train multiple policies in the source environments in an iterative manner and generalize to the variation condition with the best-performing one after the few-shot adaptation.

Open-ended Solution Discovery

Directly solving Eq. 3 or Eq. 5 requires the non-trivial pairwise constraints computation for training each policy and the parallel framework for the policy population, which limits the population size in prior works (Parker-Holder et al. 2020; Masood and Doshi-Velez 2019). Here we introduce an iterative training manner that trains only a single policy at a time and optimizes the policy π_k to be distinct from the previously discovered policies $\{\pi_i\}_{i=1}^{k-1}$, which resembles the prior work (Zhou et al. 2022). However, we concentrate on the regulated diversity objective. Formally, we optimize the following objective to train a policy π_k at each iteration:

$$\max_{\pi_k} E_{\pi_k} \left[\sum_{t=1}^T \gamma^{t-1} r_t \right] \quad (6)$$

$$\text{s.t. } E_{\pi_k} [\text{D}_{KL}[P^{\pi_k}(\cdot|f(s), f(s')) \| P^{\pi_i}(\cdot|f(s), f(s'))]] \geq \delta, \\ \forall 1 \leq i < k,$$

which converts the optimization of the whole population simultaneously to the optimization of a single policy with a simplified constraint. Furthermore, we can discover policies exhaustively through open-ended training to obtain some specific behavior pattern we desire, which is superior to prior works that fix the number of policies.

Diversity via Inverse Dynamics Disagreement

To solve the constrained optimization problem in Eq. 6, we introduce the Lagrangian multiplier method to convert the hard constraints to soft penalties:

$$\max_{\pi_k} E_{\pi_k} \left[\sum_{t=1}^T \gamma^{t-1} r_t \right] + \sum_{i=1}^{k-1} \beta_i \text{Div}(\pi_k, \pi_i), \quad (7)$$

where $\text{Div}(\pi_k, \pi_i) =$

$$E_{\pi_k} [\text{D}_{KL}[P^{\pi_k}(\cdot|f(s), f(s')) \| P^{\pi_i}(\cdot|f(s), f(s'))]],$$

where $\{\beta_i\}_{i=1}^{k-1}$ are the multipliers that can be considered as hyperparameters, and $\text{Div}(\pi_k, \pi_i)$ can be interpreted

as the inverse dynamics disagreement between two policies. Introducing the Lagrangian multipliers method to simplify the constrained optimization problem and set the multipliers as hyperparameters is widely used in the RL research (Stooke, Achiam, and Abbeel 2020; Chane-Sane, Schmid, and Laptev 2021; Peng et al. 2018). While the constraints might exhibit oscillations during training (Stooke, Achiam, and Abbeel 2020), it is acceptable since we aim to encourage diversity rather than obtain severely distinct policies.

To tractably optimize the $Div(\pi_k, \pi_i)$, we introduce an ensemble of inverse dynamics models $\{\mathcal{T}_{\phi_i}\}_{i=1}^M$ to approximate the inverse dynamics under corresponding policies. At each iteration, we train the inverse dynamics model \mathcal{T}_{ϕ_k} simultaneously by maximizing the log-likelihood:

$$L(\phi_k) = E_{\pi_k} [\log \mathcal{T}_{\phi_k}(a|f(s), f(s'))]. \quad (8)$$

To solve the optimization for policy π_k in Eq. 7, we approximate the $P^{\pi_i}(a|f(s), f(s'))$ with the an inverse dynamics model $\mathcal{T}_{\phi_i}(a|f(s), f(s'))$, and thus convert the diversity objective as follows:

$$Div_{ce}(\pi_k, \pi_i) = E_{\pi_k} [-\log \mathcal{T}_{\phi_i}(a|f(s), f(s'))], \quad (9)$$

which can be interpreted as the cross-entropy between the transitions collected by policy π_k and the inverse dynamics of π_i . We present in Appendix B.1 that the diversity objective in Eq. 9 can be approximately lower bounded by the objective in Eq. 7 with less penalty to the entropy of the policy. By introducing the novel diversity objective, we present the final objective that optimizes diversity through the transformed inverse dynamics disagreement:

$$J(\theta_k) = E_{\pi_k} \left[\sum_{t=1}^T \gamma^{t-1} r_t \right] + \frac{\alpha}{k-1} \sum_{i=1}^{k-1} Div_{ce}(\pi_k, \pi_i), \quad (10)$$

where α is a scaling hyperparameter. Since we aim to optimize the policy π_k to be distinct from each previously policies without any preference, we set the multipliers $\{\beta\}_{i=1}^{k-1}$ from Eq. 7 as $\frac{1}{k-1}$.

For implementation, we convert the diversity objective in Eq. 10 to an intrinsic reward, which trivially optimizes the objective. Herein, we define the DiR reward function as:

$$r_t^{DiR} = r_t + \frac{\alpha}{k-1} \sum_{i=1}^{k-1} [-\log \mathcal{T}_{\phi_i}(a_t|f(s_t), f(s_{t+1}))]. \quad (11)$$

We implement DiR with Proximal Policy Optimization (PPO) (Schulman et al. 2017), and we train an ensemble of policies $\{\pi_{\theta_k}\}_{k=1}^M$ and IDMs $\{\mathcal{T}_{\phi_i}\}_{i=1}^M$ iteratively. Note that we only adopt the state filtration function for the IDMs rather than the policies. The pseudo-codes of DiR and the few-shot adaptation can be found in Appendix B.2.

Connections to Prior Work

Here we denote different policies from a policy ensemble as a random variable z . Several diversity-driven approaches

maximizing the divergence between the inference action distributions of different policies on expectation (Parker-Holder et al. 2020; Derek and Isola 2021) can be formulated as maximizing $I(a; z|s) = H(a|s) - H(a|s, z)$. Additionally, the unsupervised skill discovery works focus on diversifying the state occupancy by maximizing $I(s; z) = H(s) - H(s|z)$ (Eysenbach et al. 2018) or $I(s'; z|s) = H(s'|s) - H(s'|s, z)$ (Sharma et al. 2020).

Similarly, our proposed DiR diversity objective can also be interpreted as a conditional mutual information

$$I(a; z|\bar{s}, \bar{s}') = H(a|\bar{s}, \bar{s}') - H(a|\bar{s}, \bar{s}', z), \quad (12)$$

where $\bar{s} := f(s)$ and $\bar{s}' := f(s')$. Intuitively, we encourage the output actions of different policies to be discriminable given the same partial state pair (\bar{s}, \bar{s}') . The state filtration function controls the degree of diversity. When the state filtration function is the identity function such that $f(s) = s$, there will be no further diversity as the action is relatively certain given the full state pair (s, s') . In contrast, the policies will be optimized to execute distinct actions at all time if $f(s) = \emptyset$ removes all state information. Thus, we can regulate the diverse behaviors to a specific coverage of patterns by customizing the state filtration function.

Experiment

In this section, we aim to empirically answer the following questions: (1) Can our method discover diverse behaviors? (2) Does our method discover diverse policies with regulated behavior patterns through the state filtration function? (3) Since we hypothesize DiR can obtain more feasible policies compared with baselines given the same population size, can the trained population perform better in a wide range of dynamics mismatch scenarios? Implementation details and additional results are presented in Appendix.

Experimental Settings

Environments. We adopt four continuous control tasks: Ant, Walker, Hopper, and Minitaur from Mujoco (Todorov, Erez, and Tassa 2012) and Bullet (Coumans and Bai 2016–2019), as illustrated in Appendix. C. We implement extensive test scenarios, including broken leg joints, shifted dynamics parameters, and sensor failure conditions. See Appendix C.1 for details of the environments.

Customized state filtration functions. We focus on discovering diverse locomotion patterns (e.g., walking with different legs in Ant) by designing the state filtration function for the four locomotion tasks. Thus we remove partial state information about the leg motions (e.g., joint positions) through $f(s)$ in all four environments. Full details of the state filtration functions and the original state information are described in Appendix C.2.

Baselines and implementation. We compare DiR to SMERL (Kumar et al. 2020) that trains diverse policies for generalization to environmental variations, DvD (Parker-Holder et al. 2020) that trains an ensemble of policies via the proposed divergence determinant, vanilla PPO with multiple independent policies (Multi), vanilla PPO with a single policy (PG). We set the population size as 10 in all




Figure 3: Visualization of the discovered policies in Walker. The y-axis represents the performance of discovered policies during the iterative training. We show foot contact patterns and motion illustrations of four policies in the dotted boxes. The shaded areas mark the time steps during which the respective foot (LF: left foot or RF: right foot) is in contact with the ground.

baselines, and all baselines except PG train the same number of policies. We train each policy with 2M steps for all algorithms. Each trial runs eight times with different random seeds. Refer to Appendix C.3 for implementation details.

Emergent Behavior with DiR

Since we remove the state information of one leg in Walker, we hypothesize that our method can learn various locomotion patterns (e.g., hopping, walking) with different legs. Thus, we visualize the foot contact within an episode in Walker, as shown in Figure 3. The results suggest that our method can iteratively discover diverse policies with different locomotion patterns, including hopping on both feet, incomplete one-leg hopping, and complete one-leg hopping. Furthermore, there is no significant performance degradation resulting from diversity-driven training.

Population Comparison in Training Environments

Herein we aim to compare the policies with prior diversity-driven approaches quantitatively. We adopt the population diversity, a determinant-based diversity paradigm proposed in (Parker-Holder et al. 2020), to quantify the behavior diversity of the trained policies. Here we collect 2000 states for the behavior embeddings in each environment. The results are reported in Table 1, where DiR outperforms all baseline methods concerning the diversity score. We observe that the independently trained policies in Multi can also obtain competitive diversity scores compared to DvD and SMERL which introduce extra diversity objectives, which may result from the different initialization of the policies as presented in (Jiang and Lu 2021).

To examine whether DiR can train more practical locomotion patterns by inducing the regulated diversity objective, we adopt the approach that describes behaviors of quadruped robots in prior works (Cully et al. 2015; Nilsson and Cully 2021), which computes the time proportion when the feet contact the ground within an episode. The visualization of the behavior diversity in Walker and Ant is shown in Figure 4. The results show that the behaviors of each policy trained by DiR are more consistent in multiple

episodes, and DiR can discover more distinct locomotion patterns compared with baselines.

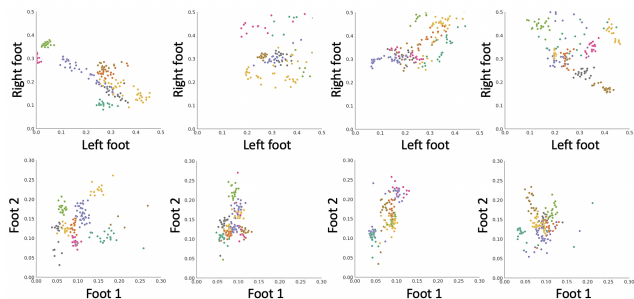


Figure 4: Visualization of behaviors through the time proportion when feet contact the ground, where the different colors indicate different policies. Different colors represent different policies from the population, and we run each policy for 20 episodes. (Top): Time proportion when two feet contact the ground within an episode in Walker. (Bottom): Time proportion when two adjacent feet contact the ground within an episode in Ant.

Furthermore, here we show the percentile performance of the population, as shown in Figure 5. The results indicate that DiR achieves competitive performance in Hopper and Minitaur while obtaining better averaged and worst-case performance in Walker and Ant, compared with baselines. We believe that regulated diversity can hinder the discovery of behavior patterns with poor performance, which results in improved worst-case performance over the policies.

Adaptation in Environment Variations

To examine whether DiR can provide adequate feasible policies which benefit the generalization, we implement various variations of the environments, including the crippled legs, the shifts of the dynamics parameters (e.g., mass), and sensor failures. See Appendix C.1 for detailed descriptions of the test conditions. For few-shot adaptation,

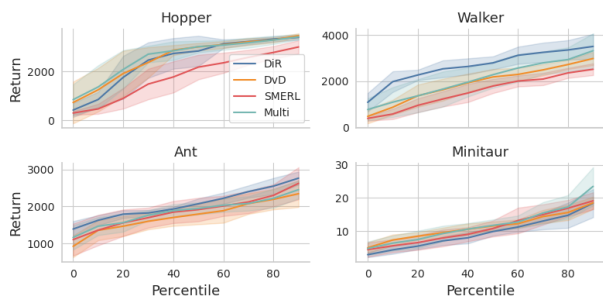


Figure 5: Percentile performance of the policy population across all four environments. The 0 percentile on the x-axis represents the worst-performing policy, while the 100 percentile represents the best-performing one.

	Ant	Hopper	Walker	Minitaur
Multi	71.6(0.4)	57.8(2.4)	76.4(0.4)	81.6(0.2)
DvD	71.6(0.5)	59.1(2.2)	76.2(0.4)	81.4(0.2)
SMERL	71.5(0.5)	63.2(2.4)	77.1(0.4)	81.7(0.3)
DiR	73.0(0.4)*	66.1(1.7)*	78.4(0.8)*	83.2(0.3)*

Table 1: Diversity scores across all environments. Asterisks indicate that the results are significantly different from all the baselines under $p < 0.05$.

we run each policy in the environment for 20 episodes and report the performance of the best-performing one.

We first evaluate the approaches under conditions with damaged body components. As the results show in Table 2, DiR outperforms baseline methods on most test conditions. In the variations of Hopper, DiR achieves comparable performance with baselines, which can result from the morphology with only one leg that limits the possible locomotion patterns. However, given the other three environments where the robots are multi-legged, DiR surpasses baselines thanks to the regulated diversity discovery. As we utilize the state filtration functions that remove the state information about the legs, DiR will focus on discovering policies that behave differently in terms of the leg motions. Thus, DiR has sufficient strategies (e.g., walking on a single leg) to handle the situations like the damaged leg. Specifically, DiR is the only approach that adapts with better locomotion patterns in the test Walker environment where a foot joint is broken. Furthermore, we remark that PG training a single policy performs significantly worse than the methods training multiple policies, which indicates that the diversity-driven approaches are simple yet effective for adaptation under dynamics variations. Importantly, we verify that the superior adaptation performance of DiR results from different policies, which further validates that the diverse behaviors benefit the extrapolation to different environments. The details of selected policies in the test environments are presented in Appendix D.1.

Furthermore, we consider the adaptation to environments

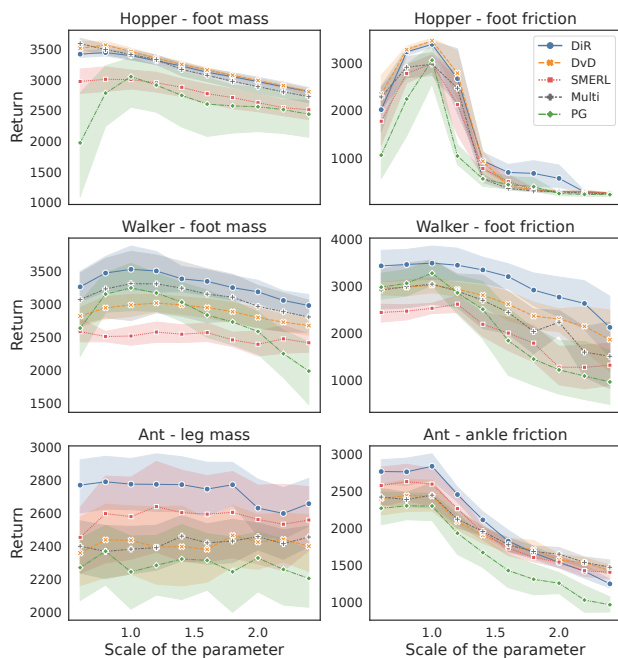


Figure 6: Adaptation performance under different levels of the dynamics parameter variations.

with shifted dynamics parameters. Here we scale the ankle friction or the leg mass of Ant, and the foot friction or the foot mass of Hopper and Walker. As the adaptation performance in Figure 6 shows, DiR outperforms baselines in Walker and Ant, while DiR produces comparable results with baselines in Hopper. Specifically, DiR provides a non-trivial improvement in *Walker - foot friction* and *Ant - leg mass* compared to all baselines, where the performance does not drop significantly when the value of foot friction gradually increases. Furthermore, we observe that DiR achieves better performance in the environment where the dynamics parameter is the same as the training environment (scale = 1), which can result from the stationary diversity-driven intrinsic rewards provided by the fixed inverse dynamics models converged in early iterations. In contrast, DvD and SMERL, whose intrinsic rewards are non-stationary due to the simultaneously trained population or discriminators, might cause unstable training dynamics and thus result in performance degradation. The phenomenon further validates the advantage of the open-ended training manner.

Finally, we implement test environments with various sensor failures where the corresponding state variables are always zeros and report the adaptation performance of the policies. The results in sensor failure conditions of Ant are shown in Table 3, where we observe that DiR is weaker than baselines when the sensors on leg 2 are defective, which might result from reliance on the state information of Leg 2 for the decision making. However, DiR outperforms baseline methods in most environments, which validates that the policies trained through DiR output actions conditioning

Environment - Damage	Multi	DvD	SMERL	DiR
Hopper - Broken leg	2972.0 ± 177.0	2798.8 ± 838.6	2193.5 ± 617.5	2587.6 ± 756.7
Hopper - Broken foot	999.5 ± 0.2	999.5 ± 0.3	999.6 ± 0.2	999.3 ± 0.4
Walker - Broken leg	2868.1 ± 315.2	2677.8 ± 268.4	2307.8 ± 221.3	3059.7 ± 199.2
Walker - Broken foot	1005.9 ± 0.5	1015.8 ± 15.7	1009.9 ± 7.0	1341.9 ± 409.6
Ant - Broken ankle	1118.7 ± 197.7	1094.8 ± 208.0	1177.0 ± 219.8	1364.5 ± 311.8
Ant - Broken hip	2222.2 ± 460.2	1998.1 ± 421.6	2148.1 ± 642.0	2534.3 ± 275.8
Minitaur - Motor failure	2.6 ± 1.3	3.0 ± 1.2	2.7 ± 1.7	3.1 ± 1.2

Table 2: Adaptation performance under the component damage.

Sensors	Multi	DvD	SMERL	DiR
Leg 1	1882 ± 258	1910 ± 196	1807 ± 249	2210 ± 408
Leg 2	1471 ± 233	1675 ± 307	1398 ± 342	1295 ± 237
Leg 3	1899 ± 323	1889 ± 183	1854 ± 269	1976 ± 260
Leg 4	1916 ± 350	2243 ± 250	2089 ± 436	2502 ± 308

Table 3: Performance under sensor failures in Ant.

on different state variables of the state information and further verify the robustness of DiR. Additional results in Walker are shown in Appendix D.2.

Related Work

In this work, we focus on the generalization across environments with various dynamics. A common approach to solve this problem is through domain randomization (Tobin et al. 2017; Peng et al. 2018), where a single policy is trained under various dynamics in simulation. Prior works have shown the effectiveness of domain randomization for the adaptation across dynamics (Rajeswaran et al. 2016; Yu et al. 2017; Akkaya et al. 2019; Shi et al. 2022; Mehta et al. 2020). Another line of work resolves the generalization through domain adaptation (Chebotar et al. 2019; Hwangbo et al. 2019; Ramos, Possas, and Fox 2019), which grounds the simulator with the collected transitions from the target domain and train the policy to be optimal under the target dynamics. Furthermore, Robust RL has shown improved transfer performance by optimizing the worst-case performance in the source environment (Pinto et al. 2017; Jiang et al. 2021; Mankowitz et al. 2019). In contrast, we resolve the generalization using an ensemble of diverse policies.

Searching for diverse solutions has been studied in Evolutionary Computation and Reinforcement Learning research. In Evolutionary Computation, Quality-Diversity (QD) is a representative type of approach that searches for diverse high-performing solutions (Cully et al. 2015; Mouret and Clune 2015; Nilsson and Cully 2021). However, the requirement of defining behavior descriptors limits QD to complicated tasks (Grillotti and Cully 2022). In Reinforcement Learning, unsupervised skill discovery has been proposed to train a latent conditional policy without environment reward (Eysenbach et al. 2018;

Sharma et al. 2020; Hartikainen et al. 2019), which can prevent the behavior from being practically feasible by ignoring the environment reward. When extrinsic rewards are also considered, several approaches have been proposed to train diverse high-performing policies (Kumar et al. 2020; Parker-Holder et al. 2020; Masood and Doshi-Velez 2019; Zhou et al. 2022; Lupu et al. 2021; Zahavy et al. 2021; Zhang, Yu, and Turk 2019). We remark that our open-ended training manner resembles the prior works that train diverse policies iteratively (Zhou et al. 2022; Zhang, Yu, and Turk 2019). However, we focus on controllable diversity through the exhaustive solution discovery different from the methods. Furthermore, several approaches resolve the generalization over various dynamics with the assistance of diverse policies (Kumar et al. 2020; Osa, Tangkaratt, and Sugiyama 2021; Kaushik, Arndt, and Kyrki 2022), same as our work. However, we take the first step to regulate the diversity for more efficient adaptation as far as we know.

Our proposed diversity optimization through the inverse dynamics disagreement also resembles Model-based RL (MBRL) (Deisenroth and Rasmussen 2019; Chua et al. 2018). The divergence between inverse dynamics models has also been proposed in the prior work for imitation learning from the observation (Yang et al. 2019). For generalization across dynamics, recent work has achieved the generalization of the dynamics model (Lee et al. 2020b; Seo et al. 2020). Unlike these works, we utilize the inverse dynamics models for diverse solution discovery.

Conclusion

In this work, we present *Diversity in Regulation* (DiR), a novel diversity-driven algorithm that learns multiple high-performing policies iteratively for adaptation under dynamics variations. The key ingredient of our method is the novel diversity objective through the inverse dynamics disagreement with the state filtration function. Specifically, we can regulate the diversity by customizing the state filtration function for desired behavior patterns. Our empirical results show that DiR can adapt to various test conditions and outperforms prior diversity-driven approaches. Overall, we believe our approach would further strengthen the understanding of diverse solution discovery and could be helpful in safe adaptation under dynamics variations which is critical for the Sim2Real problem.

Acknowledgments

This research was supported in part by Shanghai Municipal Science and Technology Major Project (No.2021SHZDZX0103), in part by Ji Hua Laboratory, Foshan, China (No.X190011TB190), in part by Science and Technology Development Center, Ministry of Education (No.2021ITA10013).

References

- Akkaya, I.; Andrychowicz, M.; Chociej, M.; Litwin, M.; McGrew, B.; Petron, A.; Paino, A.; Plappert, M.; Powell, G.; Ribas, R.; et al. 2019. Solving rubik's cube with a robot hand. *arXiv preprint arXiv:1910.07113*.
- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2016. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*.
- Andrychowicz, O. M.; Baker, B.; Chociej, M.; Jozefowicz, R.; McGrew, B.; Pachocki, J.; Petron, A.; Plappert, M.; Powell, G.; Ray, A.; et al. 2020. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1): 3–20.
- Bellman, R. 1957. A Markovian decision process. *Journal of mathematics and mechanics*, 679–684.
- Chane-Sane, E.; Schmid, C.; and Laptev, I. 2021. Goal-conditioned reinforcement learning with imagined subgoals. In *International Conference on Machine Learning*, 1430–1440. PMLR.
- Chebatar, Y.; Handa, A.; Makoviychuk, V.; Macklin, M.; Issac, J.; Ratliff, N.; and Fox, D. 2019. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. In *2019 International Conference on Robotics and Automation (ICRA)*, 8973–8979. IEEE.
- Chua, K.; Calandra, R.; McAllister, R.; and Levine, S. 2018. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *Advances in neural information processing systems*, 31.
- Coumans, E.; and Bai, Y. 2016–2019. PyBullet, a Python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>. Accessed: 2023-03-25.
- Cully, A.; Clune, J.; Tarapore, D.; and Mouret, J.-B. 2015. Robots that can adapt like animals. *Nature*, 521(7553): 503–507.
- Deisenroth, M. P.; and Rasmussen, C. E. 2019. PILCO: A model-based and data-efficient approach to policy search. In *ICML*.
- Derek, K.; and Isola, P. 2021. Adaptable Agent Populations via a Generative Model of Policies. *Advances in Neural Information Processing Systems*, 34: 3902–3913.
- Du, Y.; Watkins, O.; Darrell, T.; Abbeel, P.; and Pathak, D. 2021. Auto-tuned sim-to-real transfer. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 1290–1296. IEEE.
- Evans, B.; Thankaraj, A.; and Pinto, L. 2022. Context is Everything: Implicit Identification for Dynamics Adaptation. *arXiv preprint arXiv:2203.05549*.
- Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2018. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*.
- Grillotti, L.; and Cully, A. 2022. Unsupervised behaviour discovery with quality-diversity optimisation. *IEEE Transactions on Evolutionary Computation*.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, 1861–1870. PMLR.
- Hartikainen, K.; Geng, X.; Haarnoja, T.; and Levine, S. 2019. Dynamical distance learning for semi-supervised and unsupervised skill discovery. *arXiv preprint arXiv:1907.08225*.
- Hong, Z.-W.; Shann, T.-Y.; Su, S.-Y.; Chang, Y.-H.; Fu, T.-J.; and Lee, C.-Y. 2018. Diversity-driven exploration strategy for deep reinforcement learning. *Advances in neural information processing systems*, 31.
- Hwangbo, J.; Lee, J.; Dosovitskiy, A.; Bellicoso, D.; Tsounis, V.; Koltun, V.; and Hutter, M. 2019. Learning agile and dynamic motor skills for legged robots. *Science Robotics*, 4(26): eaau5872.
- Jaderberg, M.; Czarnecki, W. M.; Dunning, I.; Marris, L.; Lever, G.; Castaneda, A. G.; Beattie, C.; Rabinowitz, N. C.; Morcos, A. S.; Ruderman, A.; et al. 2019. Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science*, 364(6443): 859–865.
- Jiang, J.; and Lu, Z. 2021. The emergence of individuality. In *International Conference on Machine Learning*, 4992–5001. PMLR.
- Jiang, Y.; Li, C.; Dai, W.; Zou, J.; and Xiong, H. 2021. Monotonic robust policy optimization with model discrepancy. In *International Conference on Machine Learning*, 4951–4960. PMLR.
- Kaushik, R.; Arndt, K.; and Kyrki, V. 2022. SafeAPT: Safe Simulation-to-Real Robot Learning using Diverse Policies Learned in Simulation. *IEEE Robotics and Automation Letters*.
- Kumar, S.; Kumar, A.; Levine, S.; and Finn, C. 2020. One solution is not all you need: Few-shot extrapolation via structured maxent rl. *Advances in Neural Information Processing Systems*, 33: 8198–8210.
- Lee, J.; Hwangbo, J.; Wellhausen, L.; Koltun, V.; and Hutter, M. 2020a. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47): eabc5986.
- Lee, K.; Seo, Y.; Lee, S.; Lee, H.; and Shin, J. 2020b. Context-aware dynamics model for generalization in model-based reinforcement learning. In *International Conference on Machine Learning*, 5757–5766. PMLR.
- Lupu, A.; Cui, B.; Hu, H.; and Foerster, J. 2021. Trajectory diversity for zero-shot coordination. In *International Conference on Machine Learning*, 7204–7213. PMLR.
- Mankowitz, D. J.; Levine, N.; Jeong, R.; Shi, Y.; Kay, J.; Abdolmaleki, A.; Springenberg, J. T.; Mann, T.; Hester, T.; and Riedmiller, M. 2019. Robust reinforcement learning

- for continuous control with model misspecification. *arXiv preprint arXiv:1906.07516*.
- Masood, M. A.; and Doshi-Velez, F. 2019. Diversity-inducing policy gradient: Using maximum mean discrepancy to find a set of diverse policies. *arXiv preprint arXiv:1906.00088*.
- Mehta, B.; Diaz, M.; Golemo, F.; Pal, C. J.; and Paull, L. 2020. Active domain randomization. In *Conference on Robot Learning*, 1162–1176. PMLR.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.
- Mouret, J.-B.; and Clune, J. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*.
- Muratore, F.; Eilers, C.; Gienger, M.; and Peters, J. 2021. Data-efficient domain randomization with bayesian optimization. *IEEE Robotics and Automation Letters*, 6(2): 911–918.
- Nilsson, O.; and Cully, A. 2021. Policy gradient assisted map-elites. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 866–875.
- Osa, T.; Tangkaratt, V.; and Sugiyama, M. 2021. Discovering diverse solutions in deep reinforcement learning. *arXiv preprint arXiv:2103.07084*.
- Parker-Holder, J.; Pacchiano, A.; Choromanski, K. M.; and Roberts, S. J. 2020. Effective diversity in population based reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 18050–18062.
- Peng, X. B.; Andrychowicz, M.; Zaremba, W.; and Abbeel, P. 2018. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE international conference on robotics and automation (ICRA)*, 3803–3810. IEEE.
- Pinto, L.; Davidson, J.; Sukthankar, R.; and Gupta, A. 2017. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, 2817–2826. PMLR.
- Rajeswaran, A.; Ghotra, S.; Ravindran, B.; and Levine, S. 2016. Epopt: Learning robust neural network policies using model ensembles. *arXiv preprint arXiv:1610.01283*.
- Ramos, F.; Possas, R. C.; and Fox, D. 2019. Bayessim: adaptive domain randomization via probabilistic inference for robotics simulators. *arXiv preprint arXiv:1906.01728*.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International conference on machine learning*, 1889–1897. PMLR.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Seo, Y.; Lee, K.; Clavera Gilaberte, I.; Kurutach, T.; Shin, J.; and Abbeel, P. 2020. Trajectory-wise multiple choice learning for dynamics generalization in reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 12968–12979.
- Sharma, A.; Gu, S.; Levine, S.; Kumar, V.; and Hausman, K. 2020. Dynamics-Aware Unsupervised Discovery of Skills. In *International Conference on Learning Representations*.
- Shi, H.; Zhou, B.; Zeng, H.; Wang, F.; Dong, Y.; Li, J.; Wang, K.; Tian, H.; and Meng, M. Q.-H. 2022. Reinforcement learning with evolutionary trajectory generator: A general approach for quadrupedal locomotion. *IEEE Robotics and Automation Letters*, 7(2): 3085–3092.
- Stooke, A.; Achiam, J.; and Abbeel, P. 2020. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, 9133–9143. PMLR.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Tishby, N.; Pereira, F. C.; and Bialek, W. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.
- Tobin, J.; Fong, R.; Ray, A.; Schneider, J.; Zaremba, W.; and Abbeel, P. 2017. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, 23–30. IEEE.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 5026–5033. IEEE.
- Xie, Z.; Da, X.; van de Panne, M.; Babich, B.; and Garg, A. 2021. Dynamics randomization revisited: A case study for quadrupedal locomotion. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 4955–4961. IEEE.
- Yang, C.; Ma, X.; Huang, W.; Sun, F.; Liu, H.; Huang, J.; and Gan, C. 2019. Imitation learning from observations by minimizing inverse dynamics disagreement. *Advances in neural information processing systems*, 32.
- Yu, W.; Liu, C. K.; and Turk, G. 2018. Policy transfer with strategy optimization. *arXiv preprint arXiv:1810.05751*.
- Yu, W.; Tan, J.; Liu, C. K.; and Turk, G. 2017. Preparing for the unknown: Learning a universal policy with online system identification. *arXiv preprint arXiv:1702.02453*.
- Zahavy, T.; O’Donoghue, B.; Barreto, A.; Mnih, V.; Flennerhag, S.; and Singh, S. 2021. Discovering diverse nearly optimal policies with successor features. *arXiv preprint arXiv:2106.00669*.
- Zhang, Y.; Yu, W.; and Turk, G. 2019. Learning novel policies for tasks. In *International Conference on Machine Learning*, 7483–7492. PMLR.
- Zhou, Z.; Fu, W.; Zhang, B.; and Wu, Y. 2022. Continuously Discovering Novel Strategies via Reward-Switching Policy Optimization. *arXiv preprint arXiv:2204.02246*.