

Generalization through Diversity: Improving Unsupervised Environment Design

Wenjun Li, Pradeep Varakantham, Dexun Li

Singapore Management University

{wjli.2020, pradeepv, dexunli.2019}@smu.edu.sg

Abstract

Agent decision making using Reinforcement Learning (RL) heavily relies on either a model or simulator of the environment (e.g., moving in an 8x8 maze with three rooms, playing Chess on an 8x8 board). Due to this dependence, small changes in the environment (e.g., positions of obstacles in the maze, size of the board) can severely affect the effectiveness of the policy learned by the agent. To that end, existing work has proposed training RL agents on an adaptive curriculum of environments (generated automatically) to improve performance on out-of-distribution (OOD) test scenarios. Specifically, existing research has employed the potential for the agent to learn in an environment (captured using Generalized Advantage Estimation, GAE) as the key factor to select the next environment(s) to train the agent. However, such a mechanism can select similar environments (with a high potential to learn) thereby making agent training redundant on all but one of those environments. To that end, we provide a principled approach to adaptively identify diverse environments based on a novel distance measure relevant to environment design. We empirically demonstrate the versatility and effectiveness of our method in comparison to multiple leading approaches for unsupervised environment design on three distinct benchmark problems used in literature.

1 Introduction

Deep Reinforcement Learning (DRL) has had many successes in challenging tasks (e.g., Atari [Mnih *et al.*, 2015], AlphaGO [Silver *et al.*, 2016], solving Rubik’s cube [Akkaya *et al.*, 2019], chip design [Mirhoseini *et al.*, 2021]) in the last decade. However, DRL agents have been proven to be brittle and often fail to transfer well to environments only slightly different from those encountered during training [Zhang *et al.*, 2018; Cobbe *et al.*, 2019], e.g., changing size of the maze, background in a game, positions of obstacles. To achieve robust and generalizing policies, agents must be exposed to different types of environments that assist in learning different challenges [Cobbe *et al.*, 2020]. To that

end, existing research has proposed Unsupervised Environment Design (UED) that can create a curriculum (or distribution) of training scenarios that are adaptive with respect to agent policy [Mehta *et al.*, 2020; Portelas *et al.*, 2020; Wang *et al.*, 2019a; Dennis *et al.*, 2020; Jiang *et al.*, 2021b; Jiang *et al.*, 2021a]. Following the terminology in existing works, we will refer to a particular environment configuration as a *level*, e.g., an arrangement of blocks in the maze, shape/height of the obstacles in front of the agent, the geometric configuration of racing tracks, etc. In the rest of this paper, we will use levels and environments interchangeably.

1.1 Related Work on UED

Existing works on UED can be categorized along two threads. The first thread has focused on principled generation of levels and was pioneered by the Protagonist Antagonist Induced Regret Environment Design (PAIRED, [Dennis *et al.*, 2020]) algorithm. PAIRED introduced a multi-agent game between level *generator* (teacher), *antagonist* (expert student) and *protagonist* (normal student), and utilized the performance difference between *antagonist* and *protagonist* as a reward signal to guide the *generator* to adaptively create challenging training levels. The second thread has abandoned the idea of principled level generation and instead championed: (a) randomly generating levels and (b) replaying previously considered levels to deal with catastrophic forgetting by the student agent. The algorithm is referred to as Prioritized Level Replay (PLR, [Jiang *et al.*, 2021b]) and the levels are replayed based on learning potential, captured using Generalized Advantage Estimation (GAE, [Schulman *et al.*, 2015]). PLR has been empirically shown to be more scalable and also is able to achieve more robust and better “out-of-distribution” performance than PAIRED. In [Jiang *et al.*, 2021a], authors combined randomized generator and replay mechanism together and proposed PLR⁺. Empirically, PLR⁺ achieves the state-of-the-art in literature (as a method that demands no human expertise), and thus we will adopt PLR⁺ as our primary baseline in this paper. Another recent approach ACCEL [Parker-Holder *et al.*, 2022] that builds on PLR performs edits (i.e., modifications based on human expertise) on high-regret environments to learn better. Both these threads of research have significantly improved the state-of-the-art (domain randomization [Sadeghi and Levine, 2016; Tobin *et al.*, 2017] and adversarial learning [Pinto *et al.*,

2017]). However, to generalize better, it is not sufficient to train the agent only on high-regret/GAE levels. We can have multiple high-regret levels, which are very “similar” to each other and agent does not learn a lot from being trained on similar levels. Thus, levels also have to be sufficiently “different”, so that agent can gain more perspective on different challenges that it can face.

To that end, in this paper, we introduce a diversity metric in UED, which is defined based on distance between occupancy distributions associated with agent trajectories from different levels. We then provide a principled method, referred to as *Diversity Induced Prioritized Level Replay* (DIPLR) to select levels with the diversity metric to provide better generalization performance than the state-of-the-art.

1.2 Related Work on Diversity in RL

There is existing literature on diversity measurement in the field of transfer RL, e.g., [Song *et al.*, 2016; Agarwal *et al.*, 2021a; Wang *et al.*, 2019a]. However, these works conduct an exhaustive comparison of any possible states between two underlying Markov Decision Processes (MDPs), which does not take the current agent policy into account and may suffer from a large computational cost. In contrast, we use the collected trajectories induced by the policy to approximate the occupancy distribution of the different levels. Such a method can avoid massive computation and more robustly indicate the difference between two levels being explored by the same agent policy. There are a few works that have explored the diversity between policies in DRL. 1) [Hong *et al.*, 2018] adopted distance on pairwise policies to modify the loss function to enforce the agent to attempt policies different from its prior policies. As an improved version of [Hong *et al.*, 2018], [Parker-Holder *et al.*, 2020] computes the diversity on a population of policies instead of pairwise distance, so that it avoids cycling training behaviors and policy redundancy. 2) [Lupu *et al.*, 2021] proposed trajectory diversity measurement based on *Jensen-Shannon Divergence* (JSD) and included it in the loss function to maximize the difference between policies. The JSD computes the divergence based on action distribution in trajectories. In contrast to these works, our method computes both pairwise and population-wise distance for the purpose of measuring the difference between levels. Furthermore, we provide a diversity-guided method to train a robust and well-generalizing agent in the UED framework.

1.3 Contributions

Overall, our contributions are three-fold. First, we highlight the benefits of diversity in UED tasks and formally present the definition of distance between levels. Second, we employ Wasserstein distance to quantitatively measure the distance and introduce the DIPLR algorithm. Finally, we empirically demonstrate the versatility and effectiveness of DIPLR in comparison to other leading approaches on benchmark problems. Notably, we also investigate the relationship between *diversity* and *regret* (i.e., *learning potential*) and conduct an ablation study to reveal their individual effectiveness. Surprisingly, we find that diversity solely works better than regret in the model, and they complement each other to achieve the best performance when combined.

2 Background

In this section, we describe the problem of UED and also the details of the leading approaches for solving UED.

2.1 Unsupervised Environment Design, UED

The goal of UED is to train a student agent that performs well across a large set of different environments. To achieve this goal, there is a teacher agent in UED that provides a curriculum of environment parameter values to train the student agent to generalize well to unseen levels.

UED problem is formally described using an Underspecified Partially Observable Markov Decision Process (UPOMDP). It is defined using the following tuple:

$$\langle S, A, \Theta, I, O, T, R, \gamma \rangle$$

S , A and O are the set of states, actions and observations respectively. $R : S \rightarrow \mathbb{R}$ is the reward function and γ is the discount factor. The most important element in the tuple is Θ , which is the set of environment parameters. A particular parameter $\theta \in \Theta$ (can be a vector or sequence of values) defines a level and can impact the reward model, transition dynamics and the observation function, i.e., $R : S \times \Theta \rightarrow \mathbb{R}$, $T : S \times A \times \Theta \rightarrow S$ and $I : S \times \Theta \rightarrow O$. UPOMDP is underspecified, because we cannot train with all values of θ ($\in \Theta$), as Θ can be infinitely large. The goal of the student agent policy π in a UPOMDP is to maximize its discounted expected rewards for any given $\theta \in \Theta$:

$$\max_{\pi} V^{\theta}(\pi) = \max_{\pi} \mathbb{E}_{\pi} V^{\theta}(\tau) = \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^H r_t^{\theta} \cdot \gamma^t \right]$$

where r_t^{θ} is the reward obtained by π in a level with environment parameter θ at time step t . Thus, the student agent needs to be trained on a series of θ values that maximize its generalization ability on all possible levels from Θ . To that end, we employ the teacher agent. The goal of the teacher agent policy Λ is to generate a distribution over the next set of environment parameter values to train the student, i.e.,

$$\Lambda : \Pi \rightarrow \Delta(\Theta)$$

to achieve good generalization performance, where Π is the set of possible policies of the teacher.

2.2 Approaches for Solving UED

In all approaches for solving UED, the student always optimizes a policy that maximizes its value on the given θ . The different approaches for solving UED vary on the method adopted by the teacher agent. We now elaborate on the Λ employed by different approaches.

Two fundamental approaches to UED are Domain Randomization (DR) [Jakobi, 1997; Tobin *et al.*, 2017] and Minimax [Morimoto and Doya, 2005; Pinto *et al.*, 2017]. The teacher in DR simply randomizes the environment configurations regardless of the student’s policy, i.e.,

$$\Lambda^{DR}(\pi) = \mathcal{U}(\Theta)$$

where \mathcal{U} is the uniform distribution over all possible θ values. On the other hand, teacher in Minimax adversarially generates challenging environments to minimize the rewards of the student’s policy, i.e.,

$$\Lambda^{MM}(\pi) = \arg \min_{\theta} V^{\theta}(\pi)$$

The next set of approaches related to PAIRED [Dennis *et al.*, 2020; Jiang *et al.*, 2021a] rely on regret, which is defined approximately as the difference between the maximum and the mean return of students’ policy, to generate new levels:

$$reg^{\theta}(\pi) \approx \max_{\tau \sim \pi} V^{\theta}(\tau) - \mathbb{E}_{\tau \sim \pi} V^{\theta}(\tau)$$

Given this regret, teacher selects policies that maximize regret:

$$\Lambda^{MR}(\pi) = \arg \max_{\theta} reg^{\theta}(\pi)$$

From the original paper by [Dennis *et al.*, 2020], there have been multiple major improvements [Jiang *et al.*, 2021b; Parker-Holder *et al.*, 2022]:

1. New levels are generated randomly instead of using an optimizer.
2. Level to be selected at a time step is decided based on an efficient approximation of regret, referred to as *positive value loss* (a customized form of Generalized Advantage Estimation (GAE)):

$$gae^{\theta}(\pi) = \frac{1}{H} \sum_{t=0}^H \max \left(\sum_{k=t}^H (\gamma \lambda)^{k-t} \delta_k, 0 \right) \quad (1)$$

where γ and λ are the GAE and MDP discount factors respectively, H is the time horizon and δ_k is the TD-error at time step k .

3. Levels are replayed with a certain probability to ensure there is no catastrophic forgetting.
4. Edit randomly on generated levels through human-defined edits [Parker-Holder *et al.*, 2022].

We will use learning potential, regret and GAE, interchangeably to represent *positive value loss* in the rest of the paper.

2.3 Wasserstein Distance

We will propose a diversity measure that relies on distance between occupancy distributions, where we utilize Wasserstein Distance. The problem of optimal transport density between two distributions was initially proposed by [Monge, 1781] and generalized by [Kantorovich, 1942]. Wasserstein distance is preferred by the machine learning community among several commonly used divergence measures, e.g., Kullback-Leibler (KL) divergence, because it is 1) non-zero and continuously differentiable when the supports of two distributions are disjoint and 2) symmetric, i.e., $\mathcal{W}(\mathcal{P}, \mathcal{Q})$ is equal to $\mathcal{W}(\mathcal{Q}, \mathcal{P})$.

Wasserstein distance is the measurement between probability distributions defined on a metric space $M(d, \mathcal{C})$ with a cost metric $d : \mathcal{C} \times \mathcal{C} \mapsto \mathbb{R}_+$:

$$\mathcal{W}_p(\mathcal{P}, \mathcal{Q}) = \left(\inf_{\psi \in \Pi(\mathcal{P}, \mathcal{Q})} \mathbb{E}_{(x,y) \sim \psi} [d(x,y)^p] \right)^{1/p} \quad (2)$$

where $\Pi(\mathcal{P}, \mathcal{Q})$ is the set of all possible joint probability distributions between \mathcal{P} and \mathcal{Q} , and one joint distribution $\psi \in \Pi(\mathcal{P}, \mathcal{Q})$ represents a density transport plan from point x to y so as to make x follows the same probability distribution of y .

3 Approach: DIPLR

A key drawback of existing approaches for solving UED is the inherent assumption that levels with high regret (or GAE) all have high learning potential. However, if there are two levels that are very similar to each other, and they both have high regret, training the student agent on one of those levels makes training on the other unnecessary and irrelevant. Given our ultimate goal is to let student agent policy transfer well to a variety of few-shot or zero-shot levels, we utilize diversity. Diversity has gained traction in Deep Reinforcement Learning to improve the generalization performance (albeit in a different problem settings than the one in the paper) of models ([Parker-Holder *et al.*, 2020; Hong *et al.*, 2018; Lupu *et al.*, 2021]). More specifically, we provide

1. A scalable mechanism for quantitatively estimating similarity (or dissimilarity) between two given levels given the current student policy. This will then be indirectly used to generate diverse levels.
2. A teacher agent for UED that generates diverse levels and trains student agent in those diverse environments, so as to achieve strong generalization to "out-of-distribution" levels.

3.1 Estimating Similarity between Levels

The objective here is to estimate the similarity between levels given the current student agent policy. One potential option is to encode each level using one of a variety of encoding methods (e.g., Variational Autoencoder) or using the parameters associated with the level and then taking the distance between encodings of levels. However, such an approach has multiple major issues:

- Encoding a level does not account for the sequential moves the student agent will make through the level, which is of critical importance as the similarity is with respect to student policy;
- There can be stochasticity in the level that is not captured by the parameters of the level. For instance, in Bipedal-Walker, a complex yet well-parameterized UPOMDP introduced by [Wang *et al.*, 2019b; Parker-Holder *et al.*, 2022], has multiple free parameters controlling the terrain. Because of the existence of stochasticity in the level generation process, two levels could be very different while we have near-zero distance measurement given their environment parameter vectors.
- Distance between environment parameters requires normalization in each parameter dimension and is domain-specific.

Since we collect several trajectories within current levels when approximating the regret value, we can naturally get the state-action distributions induced by the current policy.

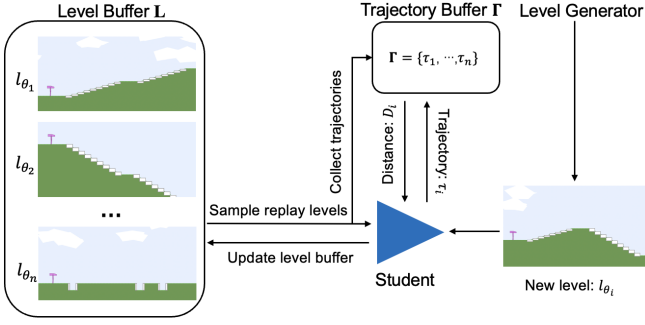


Figure 1: An Overview of DIPLR Algorithm

Therefore, we propose to evaluate *similarity on the different levels* based on the *distance between occupancy distributions of the current student policy*. The hypothesis is that if the levels are similar, then the trajectories traversed by the student agent within the level will have similar state-action distribution. Equation (3) provides the expression for state-action distribution given a policy, π and level l_{θ_1} (θ_1 represents the parameters corresponding to the level):

$$\rho_{l_{\theta_1}}^{\pi}(s, a) = (1 - \gamma) \sum_{t=0}^H \left[Pr(s_t = s, a_t = a | s_0 \sim p_0(\cdot), s_t \sim p(\cdot | s_{t-1}, a_{t-1}, \theta_1), a_t \sim \pi(\cdot | s_t)) \right] \quad (3)$$

Similarly, we can derive $\rho_{l_{\theta_2}}^{\pi}$ for level, l_{θ_2} with parameter θ_2 . Typically, KL divergence is employed to measure the distance between state-action distributions.

$$-D_{KL}(\rho_{l_{\theta_1}}^{\pi} || \rho_{l_{\theta_2}}^{\pi}) = \mathbb{E}_{(s,a) \sim \rho_{l_{\theta_1}}^{\pi}} \left[\log \frac{\rho_{l_{\theta_2}}^{\pi}}{\rho_{l_{\theta_1}}^{\pi}} \right]$$

However, KL divergence is not applicable in our setting, because of the lack of transition probabilities, the two different levels can result in two occupancy distributions with disjoint supports. More importantly, KL divergence cannot work without explicit estimates of the occupancy distributions. Therefore, we employ the Wasserstein distance described in Equation (2), which can calculate the distance between two distributions from empirical samples.

We can write the Wasserstein distance between two occupancy distributions as Equation (4):

$$\mathcal{W}(\rho_{l_{\theta_1}}^{\pi}, \rho_{l_{\theta_2}}^{\pi}) = \left(\inf_{\psi \in \Pi(\rho_{l_{\theta_1}}^{\pi}, \rho_{l_{\theta_2}}^{\pi})} \mathbb{E}_{(\phi_1, \phi_2) \sim \psi} [d(\phi_1, \phi_2)^p] \right)^{1/p} \quad (4)$$

where $\phi \in (S, A)$ is a sample from the occupancy distribution. By Equation (4), we can collect state-action samples in trajectories to compute the empirical Wasserstein distance between two levels, i.e., $\mathcal{D}(l_{\theta_1}, l_{\theta_2}) \approx \mathcal{W}(\rho_{l_{\theta_1}}^{\pi}, \rho_{l_{\theta_2}}^{\pi})$, is our empirical estimation of the Wasserstein distance between two levels.

3.2 Diversity Guided Training Agent

Now that we have a scalable way to compute the distance between levels, $\mathcal{D}(\cdot, \cdot)$, we will utilize it in the teacher to gen-

erate a diverse curriculum for the student agent. We build on the PLR training agent [Jiang *et al.*, 2021a] that only considered regret/GAE and staleness to prioritize levels for selection. Our training agent most crucially employs diversity to select levels.

Our training agent maintains a buffer of high-potential levels for training. At each iteration, we either: (a) generate a new level to be added to the buffer; or (b) sample a mini-batch of levels for the student agent to train on.

Diversity During Generating a New Level

For a new level, l_{θ} , the distance from l_{θ} to all the levels, $L = \langle l_{\theta_1}, l_{\theta_2}, \dots \rangle$ in the buffer is given by:

$$\mathcal{D}(l_{\theta}, L) = \min_k \mathcal{D}(l_{\theta}, l_{\theta_k}) \quad (5)$$

To increase diversity, we add new levels to the buffer that have the highest value of this distance (amongst all the randomly generated levels). We can also combine this distance measure with regret/GAE when taking a decision to include a level in the buffer so that different and challenging levels are added.

Diversity in Sampling Levels from Buffer to Train

To decide on levels to train on, we maintain a priority (probability of getting selected) for each level that is computed based on their distance from other levels in the buffer. The rank of a level, l_{θ_i} in terms of distance from other levels in the buffer, L is given by

$$h(\mathcal{D}_i, \langle \mathcal{D}_1, \mathcal{D}_2, \dots \rangle)$$

where $\mathcal{D}_i = \mathcal{D}(l_{\theta_i}, L \setminus l_{\theta_i})$. We then convert the rank into a probability using Equation (6),

$$P_i = \frac{1}{h(\mathcal{D}_i, \mathcal{D})^{\beta}} \quad (6)$$

where $h(\cdot, \cdot)$ is the rank transformation that can find the rank of \mathcal{D}_i in a set of values \mathcal{D} , and β is a tunable parameter.

Combining Diversity with Regret/GAE

Instead of solely considering the diversity aspect of the level buffer, we can also take learning potential into account by using regrets in Equation (1) so that we will have a level buffer that is not only filled up with diverse training levels but also challenging levels that continuously push the student. We could assign different weights to diversity and regret by letting the replay probability $P_{replay} = \rho \cdot P_D + (1 - \rho) \cdot P_R$, where P_D and P_R are the prioritization of diversity and regret respectively, and ρ is the tuning parameter.

An overview of our proposed algorithm is shown in Figure 1. When the level generator creates a new level l_{θ_i} , we collect trajectories τ_i on l_{θ_i} , and compute its distance \mathcal{D}_i to the trajectory buffer by Equation (5). If the replay probability of l_{θ_i} is greater than that of any levels in the level buffer, we insert l_{θ_i} into the level buffer and remove the level with minimum replay probability from the level buffer.

The complete procedure of DIPLR is presented in Algorithm 1. To accelerate the calculation process required by Wasserstein distance, we adopt the popular empirical Wasserstein distance solver $D(\cdot, \cdot)$ from [Flamary *et al.*, 2021]. For simplicity, we use τ to represent the state-action samples from trajectory τ in the pseudocode. To better reveal

Algorithm 1 DIPLR

Input: Level buffer size N , level generator \mathcal{G}
Initialize: student policy π_η , level buffer L , traj buffer Γ

- 1: Generate N initial levels by \mathcal{G} to populate L
- 2: Collect trajectories on each replay level in L and fill up Γ
- 3: **while** not converged **do**
- 4: Sample replay-decision, $\epsilon \sim U[0, 1]$
- 5: **if** $\epsilon \geq 0.5$ **then**
- 6: Generate a new level l_{θ_i} by \mathcal{G}
- 7: Collect trajectories τ_i on l_{θ_i} , with stop-gradient η_\perp
- 8: Compute the regret, staleness and distance for l_{θ_i}
- 9: **else**
- 10: Sample a replay level $l_{\theta_j} \in L$ according to P_{replay}
- 11: Collect trajectories τ_j on l_{θ_j} and update π_η with rewards $R(\tau_j)$
- 12: Compute the regret, staleness and distance for l_j
- 13: **end if**
- 14: Flush Γ and collect trajectories on all replay levels to fill up Γ
- 15: Update regret, staleness, and distance for l_{θ_i} or l_{θ_j}
- 16: Update L with new level l_{θ_i} if its replay probability is greater than any levels in L
- 17: Update replay probability P_{replay}
- 18: **end while**

the relationship between *diversity* and *learning potential*, we conduct an ablation study where we only adopt the diversity metric to pick levels to fill up the level buffer, i.e., set $P_{replay} = 1 \cdot P_D + 0 \cdot P_R$.

4 Experiments

In this section, we compare DIPLR to the set of leading benchmarks for solving UED: Domain Randomization, Minimax, PAIRED, PLR^\perp . We conduct experiments and empirically demonstrate the effectiveness and generality of DIPLR on three popular yet highly distinct UPOMDP domains, Minigrid, Bipedal-Walker and Car-Racing. Minigrid is a partially-observable navigation problem under discrete control with sparse rewards, while Bipedal-Walker and Car-Racing are partially-observable walking/driving under continuous control with dense rewards (we provide more details in Appendix). In each domain, we train the teacher and student agents with Proximal Policy Optimization (PPO, [Schulman *et al.*, 2017]) and we will present the zero-shot out-of-distribution (OOD) test performance of all algorithms in each domain. To make the comparison more reliable and straightforward, we adopt the recently introduced standardized DRL evaluation metrics [Agarwal *et al.*, 2021b], with which we show the aggregate inter-quartile mean (IQM) and optimality gap plots. Specifically, IQM discards the bottom and top 25% of the runs and measures the performance in the middle 50% of combined runs, and it is thus robust to outlier scores and is a better indicator of overall performance; Optimality gap captures the amount by which the algorithm fails to meet the "best performance", i.e., it assumes that a score (e.g., =1.0) is a desirable target beyond which improvements are not very important. We present the plots after normalizing the perfor-

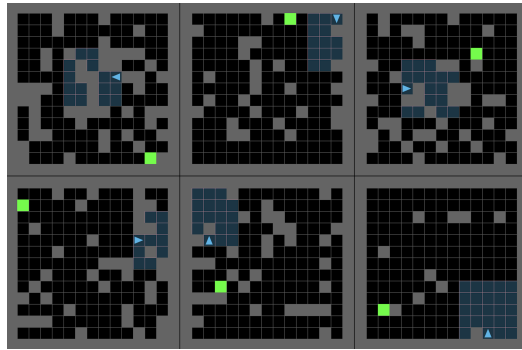


Figure 2: Examples of training environments in Minigrid Domain

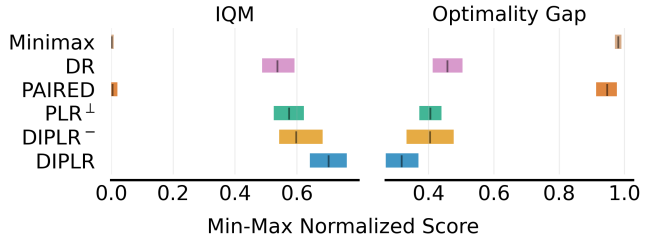


Figure 3: Aggregate zero-shot OOD test performance in Minigrid Domain across 10 independent runs. Higher IQM scores and lower optimality gap are better.

mance with a min-max range of solved-rate/returns for better visualization.

We also provide a detailed ablation analysis to demonstrate the utility of diversity alone, by providing results with and without regret (or GAE). The version with diversity alone is referred to as $DIPLR^-$ and the one which uses both diversity and regret is referred to as DIPLR.

4.1 Minigrid Domain

Introduced by [Dennis *et al.*, 2020], Minigrid is a good benchmark due to its easy interpretability and customizability. In our experimental settings, the teacher has a maximum budget of 60 steps to place small block cells and another two steps to place the start position and goal position at the end. The student agent is required to navigate to the goal cell within 256 steps and it receives rewards= $(num-steps/256)$ upon success, otherwise, it receives zero rewards. Figure 2 displays a few training examples from the domain.

We train all the student agents for 30k PPO updates (~250M steps) and evaluate their transfer capabilities on eight zero-shot OOD environments (see the first row in Figure 2) used in the literature that are crafted by humans. We summarize the empirical results of our proposed algorithm and all other benchmarks in Figure 3 and Figure 4. Here are the key observations from both the figures:

- Figure 3: DIPLR surpasses all the benchmark approaches, and surpasses the leading best approach, PLR^\perp by more than 20% with regards to IQM and Optimality Gap.

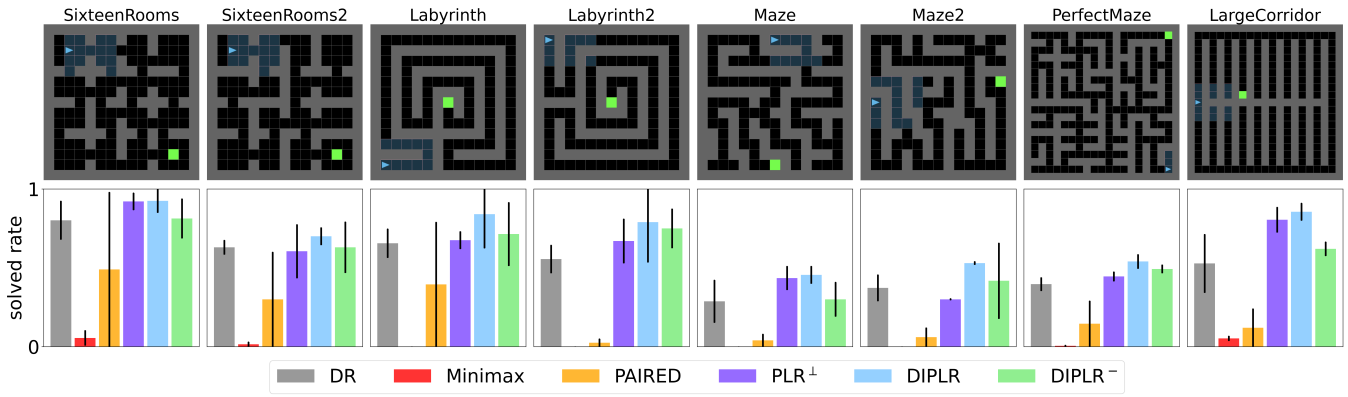


Figure 4: Zero-shot transfer performance in eight human-designed test environments. The plots are based on the median and interquartile range of solved rates across 10 independent experiments.

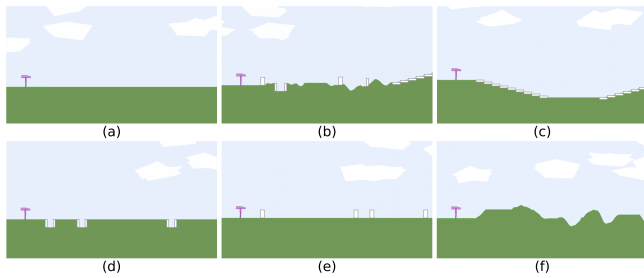


Figure 5: Six examples of test environments in Bipedal-Walker domain. (a) BipedalWalker, (b) Hardcore, (c) Stair, (d) PitGap, (e) Stump, (f) Roughness. Note that these may not be the exact test environment terrains but only extreme cases for demonstration.

- Figure 3: DIPLR⁻ performs on par or better than PLR[⊥].
- Figure 3: PAIRED performs poorly, as the long horizon (60 steps) and sparse reward task is challenging for a guided level generator. While Domain Randomization performs decently, minimax performs worse than PAIRED.
- Figure 4: With respect to zero-shot transfer results, DIPLR is obviously better than PLR[⊥] in every one of the testing scenarios. This indicates the utility of diversity.

4.2 Bipedal-Walker Domain

We also conduct a comprehensive comparison on the Bipedal-Walker domain, which is introduced in [Wang *et al.*, 2019b; Parker-Holder *et al.*, 2022] and is popular in the community as it is a well-parameterized environment with eight variables (including ground roughness, the number of stair steps, min/max range of pit gap width, min/max range of stump height, and min/max range of stair height) conditioning the terrains and is thus straightforward to interpret. The teacher in Bipedal-Walker can specify the values of the eight environment parameters, but there will still be stochasticity in the generation of a particular level. As for the student, i.e., walker agent, it should determine the torques applied on its

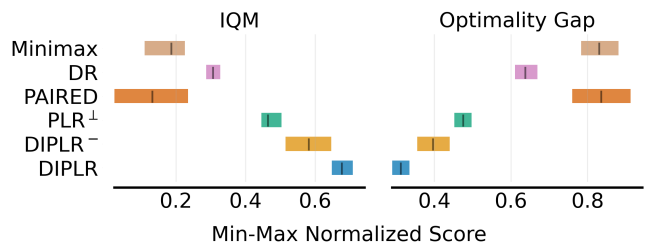


Figure 6: Aggregate test performance over 10 runs in Bipedal-Walker domain. Higher IQM scores and lower optimality gap are better

joints and is constrained by partial observability where it only knows its horizontal speed, vertical speed, angular speed, positions of joints, joints’ angular speed, etc.

Following the experiment settings in prior UED works, we train all the algorithms for 30k PPO updates (~1B steps), and then evaluate their generalization ability on six distinct test instances in Bipedal-Walker domain, i.e., BipedalWalker, Hardcore, Stair, PitGap, Stump, and Roughness, as illustrated in Figure 5. Among these, $\{BipedalWalker\}$ is the basic level that only evaluates whether the agent can walk on flat ground, $\{Stair, PitGap, Stump, Roughness\}$ challenges the walker for a particular kind of obstacles and $\{Hardcore\}$ contains a combination of all types of challenging obstacles.

To understand the evolution of transfer performance, we evaluate the student every 100 student policy updates and plot the curves in Figure 7. As is shown, our proposed method DIPLR outperforms all other benchmarks in each test environment. The ablation model DIPLR⁻ also surpasses our primary benchmark PLR[⊥] in five test environments, indicating that diversity metric alone contributes to ultimate generalization performance more than regret in continuous domains. PAIRED suffers from variance quite a bit and cannot constantly create informative levels for the student agents in such a complex domain. This is because of the nature of the multi-agent framework adopted in PAIRED, where conver-

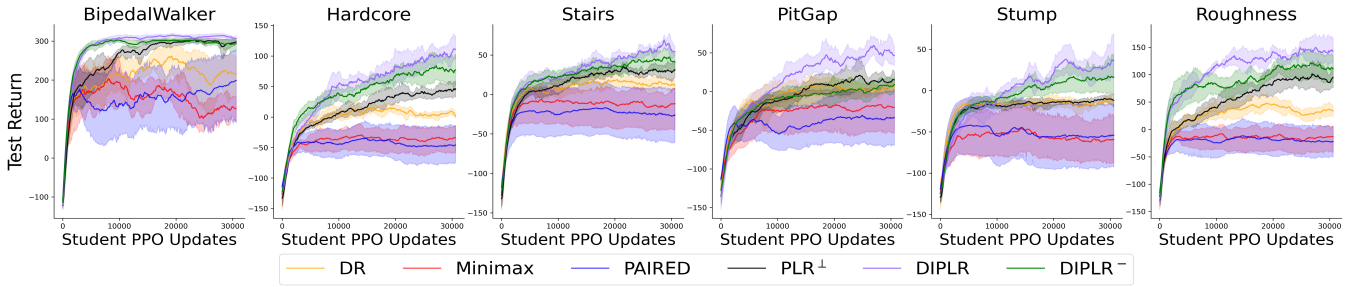


Figure 7: Performance on test environments during the training period over five independent experiments (mean and standard error).

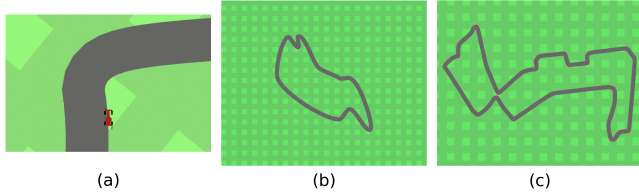


Figure 8: Examples in Car-Racing domain. (a) A zoom-in snapshot of the car-racing track; (b) a track generated by domain randomization; (c) one of the test environments, F1-Singapore track

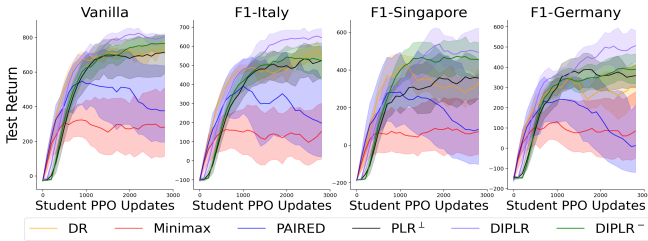


Figure 9: Zero-Shot Transfer Performance in Car-Racing domain

gence highly depends on the ability of the expert student.

After training for 30k PPO updates, we collect trained models and conduct a more rigorous evaluation based on 10 test episodes in each test environment, and present the aggregate results after min-max normalization (with range=[0, 300] on Bipedal-Walker and [0, 150] on the other five test environments) in Figure 6. Notably, our method DIPLR dominates all the benchmarks in both IQM and optimality gap. Furthermore, DIPLR⁻ performs better than PLR[⊥].

4.3 Car-Racing Domain

To validate that our method is scalable and versatile, we further implement DIPLR on the Car-Racing environment, which is introduced by [Jiang *et al.*, 2021a] and has been used in existing UED papers. In this domain, the teacher needs to create challenging racing tracks by using Bézier curves (via 12 control points) and the student drives on the track under continuous control with dense reward. A zoom-in illustration of the track is shown in Figure 8 (a). After training the student for around 3k PPO updates (~ 5.5M steps), we evaluate the student agent in four test scenarios, among which three

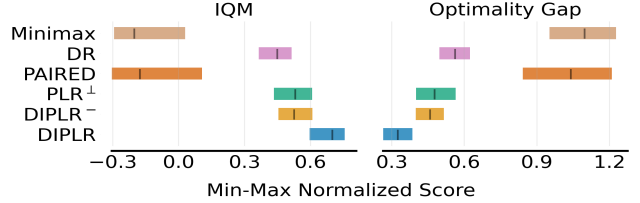


Figure 10: Aggregate test performance over 10 independent runs in Car-Racing domain.

are challenging F1-tracks existing in the real world and the other one is a basic vanilla track. Note that these tracks are absolutely OOD because they cannot be defined with just 12 control points, for example, the F1-Singapore track in Figure 8 (c) has more than 20 control points.

Figure 9 presents the complete training process of each algorithm on four test environments, with an evaluation interval of 100 PPO updates. As a relatively simpler domain, our method DIPLR and DIPLR⁻ quickly converge to the optimum while DR and PLR[⊥] achieve a near-optimal generalization on the Vanilla and F1-Italy track.

The aggregate performance after min-max normalization (with range=[200,800]) of all methods is summarized in Figure 10. Despite both PLR[⊥] and DR agents can learn well on this comparatively simple domain, DIPLR outperforms them on both IQM and optimality gap.

5 Conclusion

We provided a novel method DIPLR for unsupervised environment design (UED), where diversity in training environments is employed to improve generalization ability of a student agent. By computing the distance between levels via Wasserstein distance over occupancy distributions, we constantly expose the agent to challenging yet diverse scenarios and thus improve its generalization in zero-shot out-of-distribution test environments. In our experiments, we validated that DIPLR is capable of training robust and generalizing agents and significantly outperforms the best-performing baselines in three distinct UED domains. Moreover, we explored the relationship between *diversity* and *learning potential*, and we discover that diversity alone benefits the algorithm more than learning potential, and they complement each other to achieve the best performance when combined.

Acknowledgments

This research/project is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-RP-2020-017).

References

- [Agarwal *et al.*, 2021a] Rishabh Agarwal, Marlos C Machado, Pablo Samuel Castro, and Marc G Bellemare. Contrastive behavioral similarity embeddings for generalization in reinforcement learning. *arXiv preprint arXiv:2101.05265*, 2021.
- [Agarwal *et al.*, 2021b] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in neural information processing systems*, 34:29304–29320, 2021.
- [Akkaya *et al.*, 2019] Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- [Cobbe *et al.*, 2019] Karl Cobbe, Oleg Klimov, Chris Hesse, Taehoon Kim, and John Schulman. Quantifying generalization in reinforcement learning. In *International Conference on Machine Learning*, pages 1282–1289. PMLR, 2019.
- [Cobbe *et al.*, 2020] Karl Cobbe, Chris Hesse, Jacob Hilton, and John Schulman. Leveraging procedural generation to benchmark reinforcement learning. In *International conference on machine learning*, pages 2048–2056. PMLR, 2020.
- [Dennis *et al.*, 2020] Michael Dennis, Natasha Jaques, Eugene Vinitzky, Alexandre Bayen, Stuart Russell, Andrew Critch, and Sergey Levine. Emergent complexity and zero-shot transfer via unsupervised environment design. *Advances in neural information processing systems*, 33:13049–13061, 2020.
- [Flamary *et al.*, 2021] Remi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z Alaya, Aurelie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, et al. Pot: Python optimal transport. *J. Mach. Learn. Res.*, 22(78):1–8, 2021.
- [Hong *et al.*, 2018] Zhang-Wei Hong, Tzu-Yun Shann, Shih-Yang Su, Yi-Hsiang Chang, Tsu-Jui Fu, and Chun-Yi Lee. Diversity-driven exploration strategy for deep reinforcement learning. *Advances in neural information processing systems*, 31, 2018.
- [Jakobi, 1997] Nick Jakobi. Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive behavior*, 6(2):325–368, 1997.
- [Jiang *et al.*, 2021a] Minqi Jiang, Michael Dennis, Jack Parker-Holder, Jakob Foerster, Edward Grefenstette, and Tim Rocktäschel. Replay-guided adversarial environment design. *Advances in Neural Information Processing Systems*, 34:1884–1897, 2021.
- [Jiang *et al.*, 2021b] Minqi Jiang, Edward Grefenstette, and Tim Rocktäschel. Prioritized level replay. In *International Conference on Machine Learning*, pages 4940–4950. PMLR, 2021.
- [Kantorovich, 1942] LV Kantorovich. Kantorovich: On the transfer of masses. In *Dokl. Akad. Nauk. SSSR*, 1942.
- [Lupu *et al.*, 2021] Andrei Lupu, Brandon Cui, Hengyuan Hu, and Jakob Foerster. Trajectory diversity for zero-shot coordination. In *International Conference on Machine Learning*, pages 7204–7213. PMLR, 2021.
- [Mehta *et al.*, 2020] Bhairav Mehta, Manfred Diaz, Florian Golemo, Christopher J Pal, and Liam Paull. Active domain randomization. In *Conference on Robot Learning*, pages 1162–1176. PMLR, 2020.
- [Mirhoseini *et al.*, 2021] Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Wenjie Jiang, Ebrahim Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Azade Nazi, et al. A graph placement methodology for fast chip design. *Nature*, 594(7862):207–212, 2021.
- [Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- [Monge, 1781] Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.*, pages 666–704, 1781.
- [Morimoto and Doya, 2005] Jun Morimoto and Kenji Doya. Robust reinforcement learning. *Neural computation*, 17(2):335–359, 2005.
- [Parker-Holder *et al.*, 2020] Jack Parker-Holder, Aldo Pacchiano, Krzysztof M Choromanski, and Stephen J Roberts. Effective diversity in population based reinforcement learning. *Advances in Neural Information Processing Systems*, 33:18050–18062, 2020.
- [Parker-Holder *et al.*, 2022] Jack Parker-Holder, Minqi Jiang, Michael Dennis, Mikayel Samvelyan, Jakob Foerster, Edward Grefenstette, and Tim Rocktäschel. Evolving curricula with regret-based environment design. *arXiv preprint arXiv:2203.01302*, 2022.
- [Pinto *et al.*, 2017] Lerrel Pinto, James Davidson, and Abhinav Gupta. Supervision via competition: Robot adversaries for learning tasks. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1601–1608. IEEE, 2017.
- [Portelas *et al.*, 2020] Rémy Portelas, Cédric Colas, Katja Hofmann, and Pierre-Yves Oudeyer. Teacher algorithms for curriculum learning of deep rl in continuously parameterized environments. In *Conference on Robot Learning*, pages 835–853. PMLR, 2020.

- [Sadeghi and Levine, 2016] Fereshteh Sadeghi and Sergey Levine. Cad2rl: Real single-image flight without a single real image. *arXiv preprint arXiv:1611.04201*, 2016.
- [Schulman *et al.*, 2015] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Silver *et al.*, 2016] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [Song *et al.*, 2016] Jinhua Song, Yang Gao, Hao Wang, and Bo An. Measuring the distance between finite markov decision processes. In *Proceedings of the 2016 international conference on autonomous agents & multiagent systems*, pages 468–476, 2016.
- [Tobin *et al.*, 2017] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. 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)*, pages 23–30. IEEE, 2017.
- [Wang *et al.*, 2019a] Hao Wang, Shaokang Dong, and Ling Shao. Measuring structural similarities in finite mdps. In *IJCAI*, pages 3684–3690, 2019.
- [Wang *et al.*, 2019b] Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. Paired open-ended trailblazer (poet): Endlessly generating increasingly complex and diverse learning environments and their solutions. *arXiv preprint arXiv:1901.01753*, 2019.
- [Zhang *et al.*, 2018] Amy Zhang, Nicolas Ballas, and Joelle Pineau. A dissection of overfitting and generalization in continuous reinforcement learning. *arXiv preprint arXiv:1806.07937*, 2018.