

# Imitation from Diverse Behaviors: Wasserstein Quality Diversity Imitation Learning with Single-Step Archive Exploration

Xingrui Yu <sup>\*†</sup>  
CFAR, IHPC, Agency for Science,  
Technology and Research  
Singapore  
yu\_xingrui@cfar.a-star.edu.sg

Zhenglin Wan<sup>\*</sup>  
School of Data Science, The Chinese  
University of Hong Kong, Shenzhen  
Shenzhen, China  
vanzl3386@gmail.com

David Mark Bossens  
CFAR, IHPC, Agency for Science,  
Technology and Research  
Singapore  
david\_bossens@cfar.a-star.edu.sg

Yueming Lyu  
CFAR, IHPC, Agency for Science,  
Technology and Research  
Singapore  
lyu\_yueming@cfar.a-star.edu.sg

Qing Guo  
CFAR, IHPC, Agency for Science,  
Technology and Research  
Singapore  
guo\_qing@cfar.a-star.edu.sg

Ivor W. Tsang  
CFAR, IHPC, Agency for Science,  
Technology and Research  
Singapore  
College of Computing and Data  
Science, NTU  
Singapore  
ivor\_tsang@cfar.a-star.edu.sg

## ABSTRACT

Learning diverse and high-performance behaviors from a limited set of demonstrations is a grand challenge. Traditional imitation learning methods usually fail in this task because most of them are designed to learn one specific behavior even with multiple demonstrations. Therefore, novel techniques for *quality diversity imitation learning*, which bridges the quality diversity optimization and imitation learning methods, are needed to solve the above challenge. This work introduces Wasserstein Quality Diversity Imitation Learning (WQDIL), which 1) improves the stability of imitation learning in the quality diversity setting with latent adversarial training based on a Wasserstein Auto-Encoder (WAE), and 2) mitigates a behavior-overfitting issue using a measure-conditioned reward function with a single-step archive exploration bonus. Empirically, our method significantly outperforms state-of-the-art IL methods, achieving near-expert or beyond-expert QD performance on the challenging continuous control tasks derived from MuJoCo environments.

## KEYWORDS

Imitation Learning, Quality Diversity, Wasserstein Auto-Encoder, Wasserstein Adversarial Training, Single-Step Archive Exploration

### ACM Reference Format:

Xingrui Yu, Zhenglin Wan, David Mark Bossens, Yueming Lyu, Qing Guo, and Ivor W. Tsang. 2025. Imitation from Diverse Behaviors: Wasserstein Quality Diversity Imitation Learning with Single-Step Archive Exploration. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 16 pages.

<sup>\*</sup>Equal contribution

<sup>†</sup>Corresponding author



This work is licensed under a Creative Commons Attribution International 4.0 License.

## 1 INTRODUCTION

Imitation Learning (IL) aims to mimic an expert’s behavior by learning from the demonstrations. IL has achieved great success in many real-world applications, such as robotics [53], autonomous driving [6], and drone control [38]. However, most of the traditional imitation learning methods were designed to learn one specific behavior, even given multiple demonstrations.

Learning diverse and high-quality policies is the ultimate goal of many real-world applications like robotic locomotion tasks [4]. Recent literature has demonstrated that quality diversity reinforcement learning (QDRL) [4, 45] is a promising direction for achieving this goal. Prior methods that combine Differentiable Quality Diversity (DQD) [14] with off-policy RL achieved diverse and relatively high-performance policies. However, the performance gap between standard RL and QDRL still exists [4]. More recently, Batra et al. [4] mitigate this gap by leveraging the on-policy RL method PPO [40] with DQD. PPO estimates the gradients of diversity and performance objectives from the online collected data. Then, the estimated gradients are used by DQD methods like CMA-MAEGA [15] that maintain a single search point and move through the behavior space by filling new regions. The synergy between PPO and DQD results in a state-of-the-art QDRL method, Proximal Policy Gradient Arborecence (PPGA) [4], that achieves the ultimate goal of robotic locomotion. Generally, this kind of QDRL method finds a diverse archive of high-performing locomotion behaviors for an agent by combining PPO gradient approximations with Differentiable Quality Diversity algorithms.

However, the success of QDRL heavily relies on high-quality reward functions, which can be intractable in practice. Quality Diversity Imitation Learning (QDIL) offers a more flexible strategy for learning diverse and high-quality policies from demonstrations with diverse behaviors. In the literature, adversarial IL methods such as GAIL [23] have achieved great success in learning specific behaviors for robotic locomotion tasks. Therefore, a naive solution for QDIL is to apply adversarial IL to estimate rewards from the given demonstrations and then leverage the estimated rewards

for learning diverse and high-quality policies. Unfortunately, adversarial IL methods suffer from the **training instability** issue, which usually results in worse-than-demonstrator performance [23]. Moreover, when the demonstrations only contain a few behaviors, the rewards learned by adversarial IL techniques will be **behavior-overfitted** and unable to guide the agent to learn more diverse behaviors beyond the demonstrations.

The training instability issue and the behavior-overfitted reward issue heavily limit the adversarial QDIL to learning diverse and high-quality policies with limited demonstrations. In this work, we propose two synergic strategies to overcome these two challenging issues. The first strategy aims to stabilize the training of the reward model, and the second strategy focuses on encouraging the agent to do behavior space exploration. To develop the first strategy, we propose to stabilize the reward learning by applying Wasserstein adversarial training within the latent space of the Wasserstein Auto-Encoder (WAE) [47]. Similar to VAE [26], WAE keeps the good properties of stable training and a nice latent manifold structure while generating higher-quality images than GAN [47]. Therefore, we propose to apply WAE to enable a more stable training of reward model in adversarial QDIL. In addition, we propose latent Wasserstein adversarial training to further improve the consistency of the reward training stability. For the second strategy, we introduce a bonus that enables the agent to collect data with more diverse behaviors via single-step archive exploration. We call the resulting method Wasserstein Quality Diversity Imitation Learning (WQDIL) with Single-Step Archive Exploration (SSAE). Figure 1 illustrates the two issues of the Adversarial QDIL (i.e., training instability and behavior-overfitted reward) and the corresponding solutions (i.e. WQDIL and Single-Step Archive Exploration). The synergy between WQDIL and SSAE adherently address the two issues of adversarial QDIL, resulting in diverse and high-quality policies when learning with limited demonstrations.

We summarize our contributions as follows:

- First, we indicate the two main issues of the naive QDIL solution, i.e., the **training instability** and the **behavior-overfitted reward** issues of the adversarial QDIL approach.
- Second, we propose Wasserstein Quality Diversity Imitation Learning (WQDIL) to address the training instability issue by applying **Wasserstein adversarial training within the latent space of the Wasserstein Auto-Encoder (WAE)**.
- Third, to alleviate the behavior-overfitted reward issue, we design a **measure-conditioned reward**, which makes the reward function sensitive to the local measure space, together with a **measure-based bonus**, which encourages the agent to collect data with more diverse behaviors.

## 2 BACKGROUND AND RELATED WORK

This section provides background and related works for Imitation Learning (Section 2.1), Quality Diversity Reinforcement Learning (Section 2.2), Quality Diversity Imitation Learning (Section 2.3), Training Stability (Section 2.4), and Exploration (Section 2.5).

### 2.1 Reinforcement Learning and Imitation Learning

Reinforcement Learning (RL) searches for policies that maximize cumulative reward in an environment, typically assuming the discrete-time Markov Decision Process (MDP) formalism  $(S, A, r, P, \gamma)$ . Here,  $S$  and  $A$  are the state and action spaces,  $r(s, a)$  is the reward function,  $P(s'|s, a)$  defines state transition probabilities, and  $\gamma$  is the discount factor. The traditional RL objective is to maximize the discounted episodic return of a policy  $\mathbb{E} \left[ \sum_{k=0}^{T-1} \gamma^k r(s_k, a_k) \right]$  where  $T$  is the episode length.

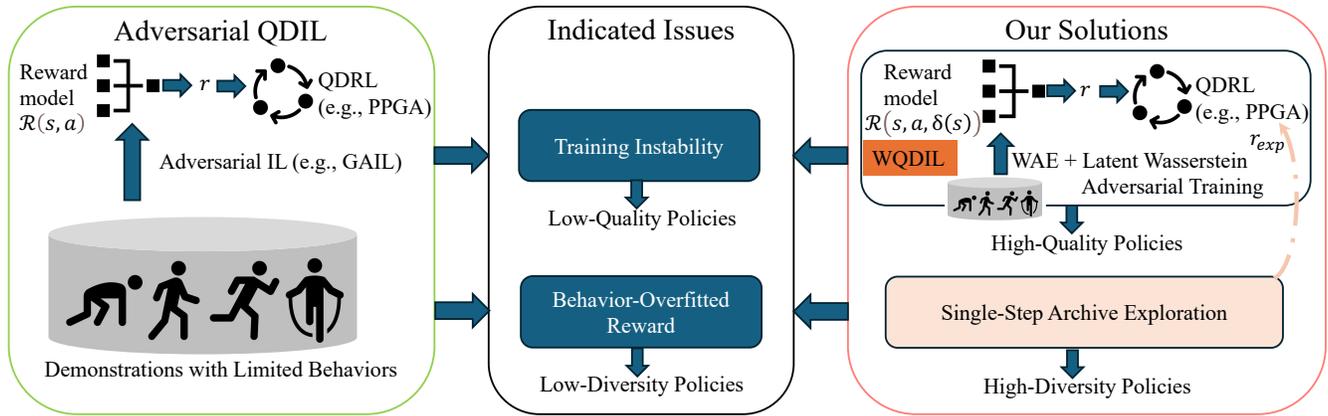
Imitation learning (IL) trains an agent to mimic expert behaviors from demonstrations [52]. Behavior Cloning (BC) uses supervised learning to imitate expert behavior but suffers from severe error accumulation [37]. Inverse Reinforcement Learning (IRL) seeks to recover a reward function from the demonstrations, and use reinforcement learning (RL) to train a policy that best mimics the expert behaviors [1]. Early IRL methods estimate rewards in the principle of maximum entropy [13, 50, 54]. Adversarial Inverse Reinforcement Learning (AIRL) [18] learns a robust reward function by training the discriminator via logistic regression to classify expert data from policy data.

**Adversarial IL** methods treat IRL as a distribution-matching problem. Generative Adversarial Imitation Learning (GAIL) [23] trains a discriminator to differentiate between the state-action distribution of the demonstrations and the state-action distribution induced by the agent’s policy, and output a reward to guide policy improvement. More recently, Primal Wasserstein Imitation Learning (PWIL) [12] introduces an offline reward function based on an upper bound of the Wasserstein distance between the expert and agent’s state-action distributions, avoiding the instability of adversarial IL methods. This paper includes max-entropy IRL (Max-EntIRL), the classic adversarial IL (i.e., GAIL) and the state-of-the-art PWIL as our baselines in QDIL with limited demonstrations.

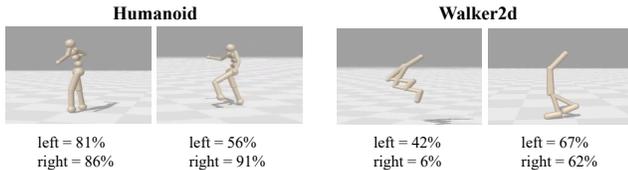
### 2.2 Quality Diversity Reinforcement Learning

**Quality Diversity Optimization.** Distinct from traditional optimization which aims to find a single solution to maximize the objective, Quality Diversity (QD) optimization aims to find a set of high-quality and diverse solutions in an  $n$ -dimensional continuous space  $\mathbb{R}^n$ . Given an objective function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  and  $k$ -dimensional measure function  $m : \mathbb{R}^n \rightarrow \mathbb{R}^k$ , the goal is to find solutions  $\theta \in \mathbb{R}^n$  for each local region in the behavior space  $B = m(\mathbb{R}^n)$ . QD algorithms discretize  $B$  into  $M$  cells, forming an archive  $\mathcal{A}$ . Formally, the objective is to find a set of solutions  $\{\theta_i\}_{i=1}^M$  which maximises  $f(\theta_i)$  for each  $i = 1, \dots, M$ . Each solution  $\theta_i$  corresponds to a cell in  $\mathcal{A}$  via its measure  $m(\theta_i)$ , forming an archive of high-quality and diverse solutions [10, 35].

Previous Quality Diversity optimization methods integrate Evolution Strategies (ES) with MAP-Elites [32], such as Covariance Matrix Adaptation MAP-Elites (CMA-ME) [16]. CMA-ME uses CMA-ES [21] as an ES algorithm generating new solutions inserted into the archive, and uses MAP-Elites to retain the highest-performing solution in each cell. CMA-ES adapts its sampling distribution based on archive improvements from offspring solutions. However, traditional ES faces low sample efficiency, especially for high-dimensional parameters such as neural networks. Differentiable



**Figure 1: Illustrations of the two issues of the Adversarial QDIL (i.e., training instability and behavior-overfitted reward) and their corresponding solutions (i.e. WQDIL and Single-Step Archive Exploration).  $\delta(s)$  means the Markovian Measure Proxy of state  $s$ , a.k.a. the single-step measure.**



**Figure 2: Illustration of diverse behaviors learned by our Quality Diversity Imitation Learning framework on Humanoid and Walker2d, where each column represents one behavior. The “left” and “right” means the proportion of time the left leg or right leg contacting the ground.**

Quality Diversity (DQD) enhances exploration and optimization by leveraging gradient information from both the objective and measure functions. The state-of-the-art DQD algorithm, CMA-MAEGA [15], utilizes CMA-ES to maintain a distribution over coefficient sets  $c_i$ , rather than directly generating solutions  $\theta$ . The objective function is  $g(\theta) = |c_0|f(\theta) + \sum_{j=1}^k c_j m_j(\theta)$ , where coefficients  $c_i$  weight the gradients of the objective  $f$  and measures  $m_j$ . CMA-MAEGA maintains a search policy and samples multiple coefficient sets after computing gradients of  $f$  and  $m_i$  to generate branched offspring solutions  $\pi_{\theta_1}, \dots, \pi_{\theta_\lambda}$ . The archive improvement from each offspring guides the update of the coefficient distribution, which is then used to refine the search policy.

Built on the QD optimization paradigm, the **Quality Diversity Reinforcement Learning (QDRL)** problem can be viewed as maximizing  $f(\theta) = \mathbb{E}_{\pi_\theta} [\sum_{k=0}^{T-1} \gamma^k r(s_k, a_k)]$  with respect to diverse  $\theta$  in a policy archive defined by measure  $m$  [11]. In QDRL, both the objective and measure are non-differentiable, requiring approximations by DQD approaches. The state-of-the-art QDRL algorithm, Proximal Policy Gradient Arborecence (PPGA), employs a vectorized PPO (VPPO) architecture to approximate the gradients of the objective and measure functions [5]. PPGA introduced the Markovian Measure Proxy (MMP), a surrogate measure function that correlates strongly with the original measure and allows gradient

approximation via policy gradient by treating it as a reward function. Specifically, MMP decomposes trajectory-based measures  $m$  into individual steps by computing:

$$m_i(\theta) = \frac{1}{T} \sum_{t=0}^T \delta_i(s_t), \quad (1)$$

where  $\delta_i(s_t)$  represents the single-step measure dependent on the state  $s_t$ . By breaking down the measure in this way, it becomes state-dependent and adheres to the Markov property. Then PPGA uses  $k+1$  parallel environments with distinct reward functions – one for the original reward and  $k$  for the surrogate measures. It approximates the gradients of both the objective and the  $k$  measure functions by comparing the policy parameters before and after multiple PPO updates. These gradients are then passed to the modified CMA-MAEGA to update the policy archive.

### 2.3 Quality Diversity Imitation Learning

*Definition 1 (Quality-Diversity Imitation Learning (QDIL)).* Given expert demonstrations  $\mathcal{D} = \{(s_i, a_i, \delta(s_i))\}_{i=1}^n$ , where  $s, a, \delta(s)$  represent state and action and the Markovian measure proxy of state  $s$ . QDIL aims to learn an archive of diverse policies  $\{\pi_{\theta_i}\}_{i=1}^M$  that collectively maximizes  $f(\theta)$  (i.e., cumulative reward) *without access to the true reward*. The archive is defined by a  $k$ -dimensional measure function  $m(\theta)$ , representing behavior patterns. After dividing the archive into  $M$  cells, the objective of QDIL is to find  $M$  solutions, each occupying one cell, to maximize:

$$\max_{\{\theta_i\}} \sum_{i=1}^M f(\theta_i). \quad (2)$$

In this paper, we focus on **Adversarial QDIL**, in which we apply adversarial IL, based on GAIL [19] in particular, to estimate rewards from the given demonstrations  $\mathcal{D}$ . This adversarial training scheme considers the policy as an adversary and an additional neural network, the discriminator  $D$ , which discriminates between samples of the policy and samples of the demonstrations. The resulting reward model is based on the discriminator, and indicates the likelihood

of the parameters based on the demonstration behaviours. This reward model is then used for the objective of QDIL.

## 2.4 Improving Stability of Adversarial IL

The methodology of adversarial IL comes from Generative Adversarial Networks (GAN) [19]. Although adversarial IL methods like GAIL have achieved great success in low-dimensional environments, they also inherit the training instability issue of GAN that limits its generalization to more diverse tasks [29]. This problem will be exacerbated in the QDIL setting due to the variety of policies being trained.

Wasserstein GAN (WGAN) [2] was proposed to stabilize the training of GANs by minimizing the Wasserstein distance instead of the Jensen-Shannon divergence between true data distribution and generated data distribution. However, the original WGAN with weight clipping can lead to undesired behaviors [20]. Gradient penalty (GP) was later proposed to address this issue, and has become a widely-used strategy for improving the stability of GAN and WGAN [20, 31, 34, 36, 39, 44]. We have applied these stabilization methods in adversarial QDIL, and conduct preliminary investigation experiments on continuous control tasks derived from MuJoCo. However, we only observed negligible effects of these methods on stabilizing adversarial QDIL.

**Wasserstein Auto-Encoder (WAE)** [47] minimizes the optimal transport cost  $W_c(P_X, P_G)$  based on the novel auto-encoder formulation. In the resulting optimization problem, the decoder tries to accurately reconstruct the encoded training examples as measured by the cost function  $f_c$ . Similar to VAE [26], the encoder tries to simultaneously achieve two conflicting goals: it tries to match the encoded distribution of training examples  $Q_Z := \mathbb{E}_{P_X}[Q(Z|X)]$  to the prior  $P_Z$  as measured by any specified divergence  $D_z(Q_Z, P_Z)$ , while making sure that the latent codes provided to the decoder are informative enough to reconstruct the encoded training examples. Using the squared cost  $f_c(x, y) = \|x - y\|_2^2$ , WAE keeps the good properties of VAEs (stable training and a nice latent manifold structure) while generating better-quality images than GAN [47]. This observation inspired us to apply WAE to improve the stability of adversarial QDIL. The resulting framework is called Wasserstein Quality Diversity Imitation Learning (WQDIL).

## 2.5 Encouraging exploration in RL and IL

Exploration bonuses have a long history in bandits and reinforcement learning as part of the “optimism in face of uncertainty” principle, and include but are not limited to visitation count based approaches. In bandits, upper confidence bound (UCB, e.g. [3]) computes a bonus based on  $\sqrt{\frac{2 \ln(t)}{n(a_t)}}$  for action  $a_t$  at time  $t$ . In tabular reinforcement learning, upper confidence reinforcement learning (UCRL) techniques (e.g. [24]) provide upper confidence bounds to the reward based on a bonus proportional to  $\sqrt{\frac{1}{n(s_t, a_t)}}$ . R-MAX [8] and E3 [25] explicitly distinguish between known and unknown states, where in unknown states R-MAX assigns the maximal reward of the MDP and E3 performs balanced wandering (i.e. taking the action with lowest visitation count for the current state). Related concepts have also been applied in deep reinforcement learning [22, 30, 41, 43]. Related to exploration bonuses, recent

techniques in imitation learning also considered curiosity-based exploration [9, 33, 51], which relates to discrepancy between the transition model’s prediction and the true transition. For instance, GIRIL [51] computes rewards offline by pretraining a reward model using a conditional VAE [42].

Our measure bonus differs from these approaches in that it is designed to explore the behavior space rather than the (state-)action space. Further, contrasting to UCRL type bonuses, due to dividing by the proportion of visits, rather than the count, our measure bonus does not shrink to zero such that underexplored behaviors will continue to receive a bonus. This is somewhat comparable in aim to the balanced wandering of E3 but differs in that it is a reward bonus rather than a policy definition (since we do not have control over the measures) and that it is maintained throughout the entire optimization process as we are interested in behavioral diversity.

## 3 WASSERSTEIN QUALITY DIVERSITY IMITATION LEARNING (WQDIL)

This section introduces our proposed framework-WQDIL. We first introduce the behavior overfitted reward problem and propose two solutions called Single-Step Archive Exploration Bonus and Measure Conditioning (Section 3.1). We then introduce Wasserstein Auto-encoders (Section 3.2), and propose a WQDIL instance WAE-WGAIL that applies latent Wasserstein adversarial training for WAE optimization (Section 3.3).

### 3.1 Mitigating behavior overfitting

Compared to the vast behavior space, expert behaviors tend to be limited in diversity. This poses a problem for learning diverse behaviors, as the reward model in the traditional IL paradigm focuses solely on shaping rewards around specific expert behaviors while neglecting other behavior patterns. This presents a challenge for balancing exploration and exploitation and makes reward functions local in that they only count a single behavior as high-performing [49]. We refer to this as the “behavior-overfitted reward” problem. This problem significantly inhibits the learning of diverse behaviors, since any RL method inherently converges to high-reward behaviors. To address this problem, we introduce a single-step archive-exploration reward bonus and measure-conditioning, which together ensure exploration across the behavior space as well as the sensitivity of the reward function to the local measure information.

*Single-Step Exploration Bonus and Measure Conditioning.* In both the bonus and the measure conditioning, we are primarily interested in single-step measures building on the Markovian Measure Proxy mentioned in Section 2.2 to make the measure information more fine-grained. This is not merely for convenience purposes but also as this will yield richer data from the demonstrations, which would otherwise be sparse (e.g. only a few episodic measures since we focus on limited-expert-demonstrations in our setting). Meanwhile, focusing on single step measure will allow exploring behaviors with similar episodic measure but different single-step measure for potentially yielding higher fitness within a particular behavioral region. This is because the episodic measure is the average (or sum) of each single-step measure and is less fine-grained.

For the exploration bonus, we establish the single-step archive  $\mathcal{A}_{single}$ , which corresponds to the state-dependent measure  $\delta(s)$ .

Similar to the behavior archive  $\mathcal{A}$ , we partition  $\mathcal{A}_{single}$  into numerous cells for discretization. Notably, instead of merely recording whether a cell is occupied, we track the **visitation count**  $n_i$  for each cell  $C_i$  in  $\mathcal{A}_{single}$ . The exploration reward bonus is formulated as

$$r_{exp}(s, a, \delta(s)) = \frac{1}{1 + p(\delta(s))}, \quad (3)$$

where

$$p(\delta(s)) = \frac{n(C(\delta(s)))}{\sum_i n(C_i)}. \quad (4)$$

In these equations,  $C_i$  denotes the  $i$ -th cell in  $\mathcal{A}_{single}$ ,  $C(\delta(s))$  represents the cell corresponding to the single-step measure proxy  $\delta(s)$ , and  $n$  signifies the visitation count. Each time a state  $s$  activates a cell in  $\mathcal{A}_{single}$ , the visitation count of that cell is incremented by one. This mechanism allows the single-step archive  $\mathcal{A}_{single}$  to be dynamically updated during training.

The exploration bonus assigns higher rewards to regions in  $\mathcal{A}_{single}$  that are less frequently visited, thereby promoting the agent to explore unseen behavior patterns. Additionally, once a region within the single-step behavior space has been sufficiently explored, the bonus decreases, facilitating the exploitation of that region to discover high-performing policies. However, note that the bonus is defined relative to the exploration of other measures such that the bonus never shrinks to zero for a particular measure. The form in Eq. 3 also avoids extreme values due to the +1 offset in the denominator. With these features together, the reward bonus can effectively mitigate the “behavior-overfitted reward” issue by always encouraging new behavior patterns, thus facilitating diverse behaviors.

In addition to the exploration bonus, the reward model is also conditioned on the measure. Specifically, using GAIL for imitation learning, the reward model is formulated as

$$\tilde{\mathcal{R}}(s, a, \delta(s)) = \max_{\pi} \min_{D} \mathbb{E}_{(s,a) \sim \mathcal{D}} [-\log D(s, a, \delta(s))] + \mathbb{E}_{(s,a) \sim \pi} [-\log(1 - D(s, a, \delta(s)))]. \quad (5)$$

This contributes to a solution to behavior overfitting: since the goal is to form behaviorally diverse policies, different behaviors require different policies to be imitated – and hence correspond to a different reward function. Combining with the exploration bonus, we obtain the combined reward model

$$\mathcal{R}(s, a, \delta(s)) = \tilde{\mathcal{R}}(s, a, \delta(s)) + r_{exp}(s, a, \delta(s)). \quad (6)$$

### 3.2 Wasserstein auto-encoders

The Wasserstein distance is a formulation that occurs within the context of optimal transport [48]. Kantorovich’s formulation of the problem is given by

$$\mathcal{W}_c(P_x, P_G) := \inf_{\Gamma \in \mathcal{P}(x \sim P_x, y \sim P_G)} \mathbb{E}_{(x,y) \sim \Gamma} [f_c(x, y)], \quad (7)$$

where  $f_c(x, y): \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_+$  is the cost function and  $\mathcal{P}(x, y)$  is the set of joint distributions of  $(x, y)$  with marginals  $P_x$  and  $P_G$  respectively. We are interested in the set  $\mathcal{X} = S \times A \times M$ , where  $M$  is the measure space.

Auto-encoders aim to represent input data in a low-dimensional latent space. Modern generative models like variational auto-encoders (VAEs) [26] and generative adversarial networks (GANs) [19] do so by minimizing a discrepancy measure between the data distribution

$P_x$  and the generative model  $P_G$ . To formulate a Wasserstein based generative model, the form in Eq. 7 is intractable; however, following [7], one can equate the Wasserstein distance to the following WAE objective [47]:

$$D_{WAE}(P_x, P_G) := \inf_{Q(z|x) \in \mathcal{Q}} \mathbb{E}_{P_x} \mathbb{E}_{Q(z|x)} [f_c(x, G(z))] + \lambda \cdot D_z(Q_z, P_z), \quad (8)$$

where  $P_z$  is the latent distribution,  $D_z$  is a divergence measure, and  $P_G$  is the decoder (i.e. generative model). We restrict our cost function to the squared Euclidian distance, i.e.  $f_c(x, y) = \|x - y\|_2^2$ , resulting in the 2-Wasserstein distance.

**WAE-GAN** [47] is an instance of WAE when choosing  $D_z(Q_z, P_z) = D_{JS}(Q_z, P_z)$  and using adversarial training to estimate  $D_z$ . Specifically, WAE-GAN uses an adversary and discriminator in the latent space  $\mathcal{Z}$  trying to separate “true” points sampled from  $P_z$  and “fake” ones sampled from  $Q_z$  [19]. In the imitation learning setting,  $P_z$  corresponds to the distribution of latent data obtained from the encoded demonstrations while  $Q_z$  corresponds to the distribution of latent data obtained from the encoded trajectory data from the policy.

It has been noted previously that WAE-GAN still suffers from instabilities due to the mode collapse property of GANs [27]. Therefore, we propose **WAE-WGAN**, which is equivalent to WAE-GAN except that it sets the divergence measure to the 1-Wasserstein distance, i.e.  $D_z(Q_z, P_z) = \mathcal{W}_1(Q_z, P_z)$ . We choose this option based on results on the improved stability during adversarial training [2]. Due to strong duality with  $p = 1$ , the Kantorovic-Rubinstein duality can be applied, i.e.

$$\mathcal{W}_1(Q_z, P_z) = \sup_{\|f\|_L \leq 1} \mathbb{E}_{z \sim Q_z} [f(z)] - \mathbb{E}_{z \sim P_z} [f(z)], \quad (9)$$

where  $f$  is a 1-Lipschitz function  $f: \mathcal{X} \rightarrow \mathbb{R}$ .

### 3.3 WQDIL: WAE Meets Adversarial QDIL

This section introduces our Wasserstein Quality Diversity Imitation Learning (WQDIL) framework, which is built on top of the Proximal Policy Gradient Arborescence (PPGA) algorithm [5]. The pseudo-code of WQDIL is given in Algorithm 1, where the blue color indicates the WQDIL’s operations that are distinct from the PPGA algorithm.

In WQDIL, we first initialize the reward model  $\mathcal{R}$  and input the expert demonstrations  $\mathcal{D}$ . Then we repeat iterations as follows. First, approximate the gradients  $\nabla f$  and  $\nabla m$  (in 1.4) using the rewards estimated from  $\mathcal{R}$  (Algorithm 2). Lines 5-15 are QD optimization options used in PPGA. In line 16, we update the search policy using the reward model  $\mathcal{R}$ . Line 17 updates the reward model  $\mathcal{R}$  using the input expert demonstrations  $\mathcal{D}$ . We now discuss different algorithms for updating the reward model  $\mathcal{R}$  in WQDIL.

**WAE-GAIL.** The first reward updating algorithm is WAE-GAIL, which is a naive adaptation of WAE-GAN in the QDIL framework. Algorithm 4 in Appendix A summarizes the details of WAE-GAIL.

**WAE-WGAIL.** Although WAE-GAIL achieves better stability and higher QD performance than GAIL in some tasks, the improvements are not consistent. To further improve WAE-GAIL, we propose to apply Wasserstein adversarial training in the latent space of WAE, which is analogous to the WAE-WGAN proposed in Section 3.2. The loss function of the auto-encoder in WAE-WGAIL is

**Algorithm 1** Wasserstein Quality Diversity Imitation Learning (WQDIL)

---

```

1: Input: Initial policy  $\theta_0$ , VPPO instance to approximate  $\nabla f$ ,
    $\nabla \mathbf{m}$  and move the search policy, number of QD iterations  $N_Q$ ,
   number of VPPO iterations to estimate the objective-measure
   functions and gradients  $N_1$ , number of VPPO iterations to move
   the search policy  $N_2$ , branching population size  $\lambda$ , and an initial
   step size for xNES  $\sigma_g$ . Initial reward model  $\mathcal{R}$ , Expert demon-
strations  $\mathcal{D}$ .
2: Initialize the search policy  $\theta_\mu = \theta_0$ . Initialize NES parameters
    $\mu, \Sigma = \sigma_g I$ 
3: for iter  $\leftarrow 1$  to  $N$  do
4:    $f, \nabla f, \mathbf{m}, \nabla \mathbf{m} \leftarrow \text{VPPO.compute\_jacobian}(\theta_\mu, \mathcal{R}, \mathbf{m}(\cdot), N_1)$ 
      $\triangleright$  approx gradients  $\nabla f$  and  $\nabla \mathbf{m}$  using the rewards estimated
     by  $\mathcal{R}$  (Algorithm 2)
5:    $\nabla f \leftarrow \text{normalize}(\nabla f), \nabla \mathbf{m} \leftarrow \text{normalize}(\nabla \mathbf{m})$ 
6:    $\_ \leftarrow \text{update\_archive}(\theta_\mu, f, \mathbf{m})$ 
7:   for  $i \leftarrow 1$  to  $\lambda$  do // branching solutions
8:      $c \sim \mathcal{N}(\mu, \Sigma)$  // sample gradient coefficients
9:      $\nabla_i \leftarrow c_0 \nabla f + \sum_{j=1}^k c_j \nabla m_j$ 
10:     $\theta'_i \leftarrow \theta_\mu + \nabla_i$ 
11:     $f', *, \mathbf{m}', * \leftarrow \text{rollout}(\theta'_i, \mathcal{R})$ 
12:     $\Delta_i \leftarrow \text{update\_archive}(\theta'_i, f', \mathbf{m}')$   $\triangleright$  get archive improve-
     ment of each solution using Algorithm 2 of PPGA paper.
13:   Rank gradient coefficients  $\nabla_i$  by archive improvement  $\Delta_i$ 
14:   Adapt xNES parameters  $\mu = \mu', \Sigma = \Sigma'$  based on improve-
     ment ranking  $\Delta_i$ 
15:    $f'(\theta_\mu) = c_{\mu,0} f + \sum_{j=1}^k c_{\mu,j} m_j$ , where  $c_\mu = \mu'$ 
16:    $\theta'_\mu = \text{VPPO.train}(\theta_\mu, f', \mathbf{m}', N_2, \mathcal{R})$   $\triangleright$  walk search policy us-
     ing reward model  $\mathcal{R}$ 
17:    $\mathcal{R}.\text{update}(\mathcal{D}, \theta'_\mu)$   $\triangleright$  update reward model using Algorithm 4
     (Appendix A) for WAE-GAIL, Algorithm 5 (Appendix A) for
     WAE-WGAIL, or Algorithm 6 for mCWAE-WGAIL.
18:   if there is no change in the archive then
19:     Restart xNES with  $\mu = 0, \Sigma = \sigma_g I$ 
20:     Set  $\theta_\mu$  to a randomly selected existing cell  $\theta_i$  from the
     archive
21: Output: Updated policy archive.

```

---

**Algorithm 2** Reward Estimation

---

```

1: Initialize: Reward model  $\mathcal{R}$ 
2: Method: Reward estimation in Algorithm 1 (step 4)
3: def get_episode_reward(episode, current archive  $\mathcal{A}$ ):
4:    $s_1, a_1, \delta(s_1), s_2, a_2, \delta(s_2), \dots, s_k, a_k, \delta(s_k) \leftarrow \text{episode}()$ 
5:    $r_1, r_2, \dots, r_k \leftarrow \mathcal{R}(s, a, \delta(s))$   $\triangleright$  batch reward (see Eq.5)
6:   compute exploration bonuses  $r_{1,exp}, r_{2,exp}, \dots, r_{k,exp}$   $\triangleright$  (see
   Eq.3)
7:   For  $i = 1, \dots, k$ 
8:      $r_i \leftarrow r_i + r_{i,exp}$   $\triangleright$  calculate total reward for each step
9:   return  $r_1, r_2, \dots, r_k$ 

```

---

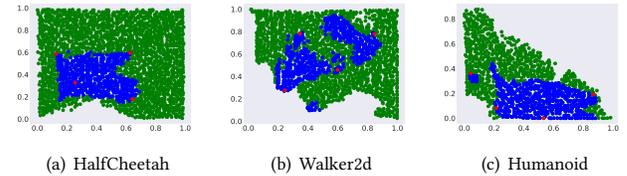
based on the reconstruction loss for expert demonstrations, the reconstruction loss for policy data, and the Wasserstein distance

between latent variables of expert demonstrations and policy data:

$$\begin{aligned}
\mathcal{L}(\phi, \psi) &= \mathcal{L}_{\text{recon}}^e + \mathcal{L}_{\text{recon}}^\pi + \lambda D_z(z^e, z^\pi) \\
&= \mathbb{E}_{x^e \sim \mathcal{D}, z^e \sim Q_\phi(z^e | x^e)} [f_c(x^e, G_\psi(z^e))] \\
&\quad + \mathbb{E}_{x^\pi \sim \pi, z^\pi \sim Q_\phi(z^\pi | x^\pi)} [f_c(x^\pi, G_\psi(z^\pi))] \\
&\quad + \lambda \mathbb{E}_{x^e \sim \mathcal{D}, x^\pi \sim \pi} [D_{\mathcal{W}}(G_\psi(x^e)) - D_{\mathcal{W}}(G_\psi(x^\pi))], \quad (10)
\end{aligned}$$

where  $\phi$  and  $\psi$  are the parameters of the encoder and decoder, respectively,  $x = (s, a)$  is the state-action pair,  $Q_\phi$  is the encoder,  $G_\psi$  is the decoder, and  $D_{\mathcal{W}}$  is the Wasserstein discriminator. Algorithm 5 in Appendix A summarizes the details of WAE-WGAIL.

**mCWAE-WGAIL.** Compared to GAIL, WAE-WGAIL has significantly better stability and consistency across different tasks. However, the latent space of WAE-WGAIL is learned from the state-action pairs. Since the demonstrations only contain very limited behaviors, the learned latent space is not adaptive to more diverse behaviors that the agent may encounter. To improve the adaptivity of the latent space and the reward model  $\mathcal{R}$ , we propose measure-conditioned WAE-WGAIL (mCWAE-WGAIL) that learns latent space by reconstructing state-action-measure triplet, i.e. formulating  $x = (s, a, \delta(s))$  in Eq. 10. Algorithm 6 summarizes the details of mCWAE-WGAIL, our proposed algorithm for WQDIL.



**Figure 3: Visualization of the behavior space. The x and y axes are the proportions of time leg 1 and 2, respectively, touch the ground. Green indicates the full expert behavior space, blue indicates the selected top-500 elites, and red indicates the selected demonstrators.**

## 4 EXPERIMENTS

### 4.1 Experiment Setup

We evaluate our framework on three popular MuJoCo [46] environments: Halfcheetah, Humanoid, and Walker2d. The measure function is the number of times a leg contacts with the ground, divided by the trajectory length. We implement all the methods based on the PPGA implementation, which utilizes the Brax simulator [17] with QDax wrappers for measure calculation [28]. All experiments were conducted on A40 GPUs for three runs per task.

### 4.2 Demonstrations

We use a policy archive obtained by PPGA to generate expert demonstrations. To follow real-world scenario with limited demonstrations, we first sample the top 500 high-performance elites from the archive as candidate pool and then select 4 policies from these candidates such that they are as diverse as possible. We generate 1 demonstration for each selected policy, resulting in the final 4

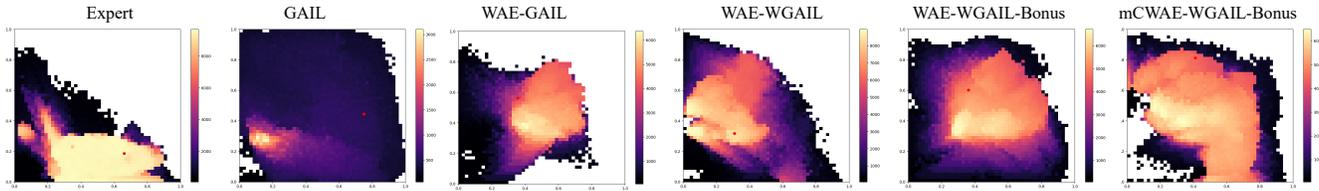


Figure 4: Visualization of the policy archive of Expert, GAIL, WAE-GAIL, WAE-WGAIL, WAE-WGAIL-Bonus and mCWAE-WGAIL-Bonus on Humanoid. The color indicates the cumulative rewards of best performing policy in the archive cells.

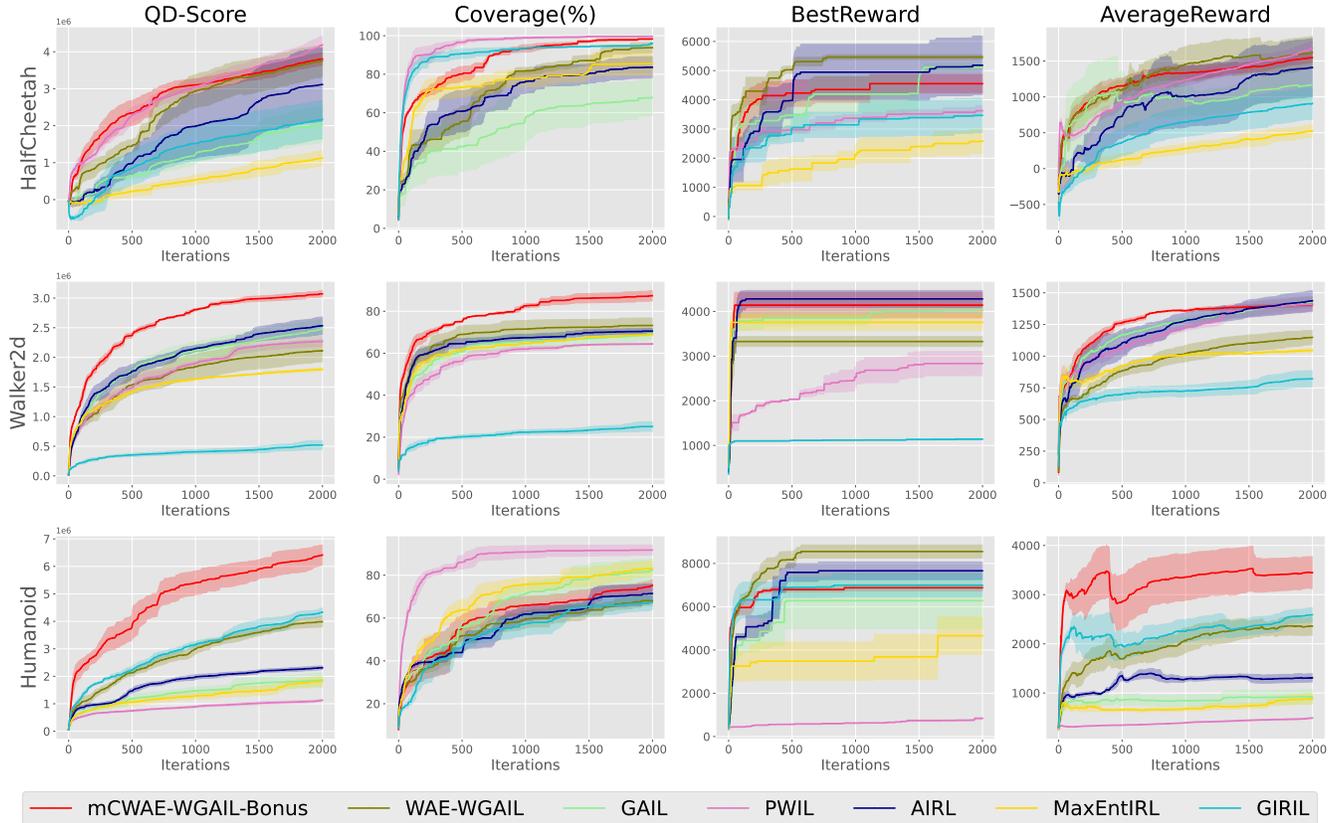


Figure 5: Learning curve comparison of our mCWAE-WGAIL-Bonus against WAE-WGAIL and state-of-the-art IL methods. The curves (and shaded areas) represent the means (and standard deviations) of the algorithms. Columns indicate the metric (QD-Score, Coverage, Best Reward, and Average Reward) while the rows represent the different benchmarks (Halfcheetah, Walker2d, and Humanoid).

demonstrations for QDIL. This process results in 4 diverse demonstrations (episodes) per environment. Figure 3 visualizes the demonstrators in the policy archive, and Table 4 in Appendix C provides the statistical properties of the demonstrations.

### 4.3 Performance

To validate the effectiveness of WQIL with Single-Step Archive Exploration, we compare our mCVAE-WAIL-Bonus against state-of-the-art IL methods within QDIL framework. These methods includes our WAE-WGAIL (Algorithm 5, Appendix A), GAIL [23],

PWIL [12], AIRL [18], MaxEntIRL [54] and GIRIL [51]. We apply a gradient penalty to all the methods that require the online update of the reward model. Appendix B summarizes the hyperparameters.

We compare the performance under four QD metrics [5]: 1) **QD-Score**, the sum of scores of all nonempty cells in the archive. QD-score is the most important metric in QDIL as it aligns with the objective of QDIL as in equation (2); 2) **Coverage**, the percentage of nonempty cells, indicating the algorithm’s ability to discover diverse behaviors; 3) **Best Reward**, the highest score found by the algorithm; and 4) **Average Reward**, the mean score of all

nonempty cells, reflecting the ability to discover both diverse and high-performing policies. We use the true reward functions to calculate these metrics.

The heatmaps in Figure 4 visualize the performance across the policy archive for Expert, GAIL, WAE-GAIL, WAE-WGAIL, WAE-WGAIL-Bonus and mCWAE-WGAIL-Bonus on Humanoid. Even using true rewards, the expert explores less than half of the archive in the Humanoid environment. The high-performance policies are more restricted, covering around a quarter of the full archive. GAIL nearly explores the full archive space, however, most of the learned policies are low-performance (in dark color). Due to the training instability and behavior-overfitting issues, GAIL only learns high-performance policies within a small region of the behavior space. WAE-GAIL learns more high-performance policies than GAIL since it maintains better stability than GAIL. However, the explored region is still limited. By applying latent Wasserstein adversarial training, WAE-WGAIL further enlarges the explored archive region and coverage of high-performance policies. WAE-WGAIL-Bonus updates the archive using the rewards learned by WAE-WGAIL and the estimated single-step archive exploration bonus. The explored region of WAE-WGAIL-Bonus is clearly larger than WAE-WGAIL, which we attribute to the exploration bonus yielding improved coverage of the measure space. Thanks to the measure conditioning, the explored archive space of mCWAE-WGAIL-Bonus is slightly smaller but is more focused in the sense that it yields the largest set of high-performing policies among the five QDIL methods.

Figure 5 shows the learning curve comparison of QD performance for different QDIL methods. Among the four metrics, QD-Score is the most important because it reflects both high diversity and high quality of performance. In HalfCheetah, our method mCWAE-WGAIL-Bonus achieves one of the highest QD-Score with much better sample efficiency than the other counterparts (i.e., achieving the same QD-Score with much fewer iterations). In Walker2d and Humanoid environments, our method achieves a significantly higher QD-Score than the other methods. PWIL achieves a slightly higher QD-Score than our method in HalfCheetah, however, it performs the worst in Humanoid. MaxEntIRL performs worse than GAIL, while AIRL outperforms GAIL slightly. Pretraining from limited demonstrations, GIRL performs similar to GAIL in HalfCheetah and similar to WAE-WGAIL in Humanoid but much worse in Walker2d. For Coverage, our mCWAE-WGAIL-Bonus is the best in HalfCheetah and Walker2d. Although our coverage in Humanoid is not the best, we significantly outperform the baselines in terms of AverageReward. The quantitative results are summarized in Table 5 of Appendix D.1.

#### 4.4 Ablations

We conduct two ablations to study the effects of each component of our method mCWAE-WGAIL-Bonus. The first ablation in Section 4.4.1 studies the effects of latent Wasserstein adversarial training in the QDIL framework. On the other hand, Section 4.4.2 focuses on studying the effects of single-step archive exploration bonus and measure conditioning.

**4.4.1 Ablation on Latent Wasserstein Adversarial Training.** To study the effect of latent Wasserstein adversarial training, we additionally introduce mCWAE-GAIL-Bonus, which reinforces WAE-GAIL

with measure conditioning and a single-step archive exploration bonus. Table 7 in Appendix D.2 compares the QD performance of our mCWAE-WGAIL-Bonus with mCWAE-GAIL-Bonus and WAE-WGAIL. As we can see, applying Wasserstein adversarial training improves the QD-Score of WAE-GAIL and mCWAE-GAIL-Bonus by 27.5% and 23.7% respectively on HalfCheetah. In Walker2d, the latent Wasserstein adversarial training enhances the QD-Score of WAE-GAIL and mCWAE-GAIL-Bonus by 74.3% and 62.4%, which is significant. Moreover in Humanoid, mCWAE-WGAIL-Bonus achieves 2x QD-Score of mCWAE-GAIL-Bonus, impressively outperforming the expert (i.e., PPGA-trueReward) by 12%. The learning curves are shown in Figure 6 in Appendix D.2.

**4.4.2 Ablations on Single-Step Archive Exploration and Measure Conditioning.** Table 8 in of Appendix D.2 compares the effects of single-step archive exploration bonus and measure conditioning. Measure conditioning improves the QD-Score for WAE-WGAIL-Bonus by 42.8%. In Walker2d, the contributions of measure conditioning to WAE-WGAIL and WAE-WGAIL-Bonus are 12.3% and 20.3% respectively. In Humanoid, mCWAE-WGAIL-Bonus outperforms WAE-WGAIL-Bonus by 28.7%. Now, we study the effect of the single-step archive exploration bonus. The bonus does not show much improvement on HalfCheetah. Comparing mCWAE-WGAIL-Bonus and mCWAE-WGAIL, we found that the contributions of bonus on the QD-Score improvement are 29.5% in Walker2d and 80% in Humanoid. Comparing WAE-WGAIL-Bonus and WAE-WGAIL, the contributions of the bonus on the QD-Score improvement are 20.8% in Walker2d and 24.8% in Humanoid. We illustrate the corresponding learning curves of this ablation study in Figure 7 of Appendix D.2.

## 5 CONCLUSION

This paper focuses on addressing two key issues of adversarial QDIL, i.e., the training instability of adversarial training and the behavior-overfitting in the design of the reward function. We propose WQDIL and a measure-conditioned reward function with Single-Step Archive Exploration to address these two issues respectively. Through two thorough ablation studies, we confirm the significance of each component (latent Wasserstein training, measure conditioning, and single-step exploration bonus) and find that latent Wasserstein adversarial training contributes most significantly to the QD-Score. The collaboration between these components contributes to the superior QD performance of our method on three continuous control tasks. Our method enables the agent to stably learn diverse and high-quality behaviors from limited demonstrations.

## ACKNOWLEDGMENTS

This project is supported by the National Research Foundation, Singapore and Infocomm Media Development Authority under its Trust Tech Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of National Research Foundation, Singapore and Infocomm Media Development Authority. Zhenglin Wan contributed to this publication whilst serving his internship with A\*STAR CFAR under the Foreign Student Attachment scheme.

## REFERENCES

- [1] Pieter Abbeel and Andrew Y Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first International Conference on Machine Learning*.
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *International conference on machine learning*. PMLR, 214–223.
- [3] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine Learning* 47 (2002), 235–256. <https://doi.org/10.1023/A:1013689704352>
- [4] Sumeet Batra, Bryon Tjanaka, Matthew C Fontaine, Aleksei Petrenko, Stefanos Nikolaidis, and Gaurav Sukhatme. 2023. Proximal Policy Gradient Arborecence for Quality Diversity Reinforcement Learning. *arXiv preprint arXiv:2305.13795* (2023).
- [5] Sumeet Batra, Bryon Tjanaka, Matthew C Fontaine, Aleksei Petrenko, Stefanos Nikolaidis, and Gaurav Sukhatme. 2023. Proximal policy gradient arborecence for quality diversity reinforcement learning. *arXiv preprint arXiv:2305.13795* (2023).
- [6] Mariusz Bojarski. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316* (2016).
- [7] O. Bousquet, S. Gelly, I. Tolstikhin, C. J. Simon-Gabriel, and B. Schölkopf. 2017. From optimal transport to generative modeling: the VEGAN cookbook. *arXiv:arXiv:1705.07642*
- [8] Ronen I Brafman and Moshe Tennenholtz. 2002. R-MAX – A General Polynomial Time Algorithm for Near-Optimal Reinforcement Learning. *Journal of Machine Learning Research* 3 (2002), 213–231. <https://doi.org/10.1162/153244303765208377>
- [9] Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A Efros. 2019. Large-scale study of curiosity-driven learning. In *International Conference on Learning Representations*.
- [10] Konstantinos Chatzilygeroudis, Antoine Cully, Vassilis Vassiliades, and Jean-Baptiste Mouret. 2021. Quality-diversity optimization: a novel branch of stochastic optimization. In *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*. Springer, 109–135.
- [11] Geoffrey Cideron, Thomas Pierrot, Nicolas Perrin, Karim Beguir, and Olivier Sigaud. 2020. Qd-rl: Efficient mixing of quality and diversity in reinforcement learning. *arXiv preprint arXiv:2006.08505* (2020), 28–73.
- [12] Robert Dadashi, Léonard Hussonot, Matthieu Geist, and Olivier Pietquin. 2021. Primal wasserstein imitation learning. In *International Conference on Learning Representations*.
- [13] Chelsea Finn, Sergey Levine, and Pieter Abbeel. 2016. Guided cost learning: Deep inverse optimal control via policy optimization. In *International Conference on Machine Learning*. PMLR, 49–58.
- [14] Matthew Fontaine and Stefanos Nikolaidis. 2021. Differentiable quality diversity. *Advances in Neural Information Processing Systems* 34 (2021), 10040–10052.
- [15] Matthew Fontaine and Stefanos Nikolaidis. 2023. Covariance matrix adaptation map-annealing. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 456–465.
- [16] Matthew C Fontaine, Julian Togelius, Stefanos Nikolaidis, and Amy K Hoover. 2020. Covariance matrix adaptation for the rapid illumination of behavior space. In *Proceedings of the 2020 genetic and evolutionary computation conference*. 94–102.
- [17] C Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. 2021. Brax—a differentiable physics engine for large scale rigid body simulation, 2021. URL <http://github.com/google/brax> 6 (2021).
- [18] Justin Fu, Katie Luo, and Sergey Levine. 2017. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248* (2017).
- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*. 2672–2680.
- [20] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. 2017. Improved training of wasserstein gans. *Advances in neural information processing systems* 30 (2017).
- [21] Nikolaus Hansen and Andreas Ostermeier. 2001. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation* 9, 2 (2001), 159–195.
- [22] Mikael Henaff, Roberta Raileanu, Minqi Jiang, and Tim Rocktäschel. 2022. Exploration via Elliptical Episodic Bonuses. In *Advances in Neural Information Processing Systems (NeurIPS 2022)*.
- [23] Jonathan Ho and Stefano Ermon. 2016. Generative adversarial imitation learning. *Advances in neural information processing systems* 29 (2016).
- [24] Thomas Jaksch, Ronald Ortner, and Peter Auer. 2010. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research* 11 (2010), 1563–1600.
- [25] Michael Kearns and Satinder Singh. 2002. Near-optimal reinforcement learning in polynomial time. *Machine Learning* 49, 2-3 (2002), 209–232. <https://doi.org/10.1023/A:1017984413808>
- [26] D. P. Kingma and M. Welling. 2014. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*.
- [27] Soheil Kolouri, Phillip E. Pope, Charles E. Martin, and Gustavo K. Rohde. 2019. Sliced Wasserstein Auto-Encoders. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=H1xaJn05FQ>
- [28] Bryan Lim, Maxime Allard, Luca Grillotti, and Antoine Cully. 2022. Accelerated quality-diversity for robotics through massive parallelism. In *ICLR Workshop on Agent Learning in Open-Endedness*.
- [29] Tianjiao Luo, Tim Pearce, Huayu Chen, Jianfei Chen, and Jun Zhu. 2024. C-GAIL: Stabilizing Generative Adversarial Imitation Learning with Control Theory. *arXiv preprint arXiv:2402.16349* (2024).
- [30] Marlos C. Machado, Marc G. Bellemare, and Michael Bowling. 2020. Count-based exploration with the successor representation. *AAAI 2020 - 34th AAAI Conference on Artificial Intelligence (2020)*, 5123–5133. arXiv:1807.11622
- [31] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. 2018. Which training methods for GANs do actually converge?. In *International conference on machine learning*. PMLR, 3481–3490.
- [32] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
- [33] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. 2017. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning*. PMLR, 2778–2787.
- [34] Henning Petzka, Asja Fischer, and Denis Lukovnikov. 2017. On the regularization of wasserstein gans. *arXiv preprint arXiv:1709.08894* (2017).
- [35] Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. 2016. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI* 3 (2016), 202845.
- [36] Guo-Jun Qi. 2020. Loss-sensitive generative adversarial networks on lipschitz densities. *International Journal of Computer Vision* 128, 5 (2020), 1118–1140.
- [37] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 627–635.
- [38] Stéphane Ross, Narek Melik-Barkhudarov, Kumar Shaurya Shankar, Andreas Wendel, Debadepta Dey, J Andrew Bagnell, and Martial Hebert. 2013. Learning monocular reactive uav control in cluttered natural environments. In *2013 IEEE international conference on robotics and automation*. IEEE, 1765–1772.
- [39] Kevin Roth, Aurelien Lucchi, Sebastian Nowozin, and Thomas Hofmann. 2017. Stabilizing training of generative adversarial networks through regularization. *Advances in neural information processing systems* 30 (2017).
- [40] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [41] Younggyo Seo, Lili Chen, Jinwoo Shin, Honglak Lee, Pieter Abbeel, and Kimin Lee. 2021. State Entropy Maximization with Random Encoders for Efficient Exploration. *Proceedings of Machine Learning Research* 139 (2021), 9443–9454.
- [42] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems (NeurIPS 2015)*.
- [43] Haoran Tang, Rein Houthoofd, Davis Foote, Adam Stooke, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. 2017. #Exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS 2017)*. 2754–2763.
- [44] Hoang Thanh-Tung, Truyen Tran, and Svetha Venkatesh. 2019. Improving generalization and stability of generative adversarial networks. *arXiv preprint arXiv:1902.03984* (2019).
- [45] Bryon Tjanaka, Matthew C Fontaine, Julian Togelius, and Stefanos Nikolaidis. 2022. Approximating gradients for differentiable quality diversity in reinforcement learning. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 1102–1111.
- [46] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 5026–5033.
- [47] Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schölkopf. 2018. Wasserstein auto-encoders. In *International Conference on Learning Representations*.
- [48] C. Villani. 2003. *Topics in Optimal Transportation*. AMS Graduate Studies in Mathematics.
- [49] Zhenglin Wan, Xingrui Yu, David Mark Bossens, Yueming Lyu, Qing Guo, Flint Xiaofeng Fan, and Ivor Tsang. 2024. Quality Diversity Imitation Learning. (2024). arXiv:2410.06151 [cs.LG] <https://arxiv.org/abs/2410.06151>
- [50] Markus Wulfmeier, Peter Ondruska, and Ingmar Posner. 2015. Maximum entropy deep inverse reinforcement learning. *arXiv preprint arXiv:1507.04888* (2015).
- [51] Xingrui Yu, Yueming Lyu, and Ivor Tsang. 2020. Intrinsic reward driven imitation learning via generative model. In *International Conference on Machine Learning*. PMLR, 10925–10935.
- [52] Maryam Zare, Parham M Kebria, Abbas Khosravi, and Saeid Nahavandi. 2024. A survey of imitation learning: Algorithms, recent developments, and challenges. *IEEE Transactions on Cybernetics* (2024).

- [53] Yuke Zhu, Ziyu Wang, Josh Merel, Andrei Rusu, Tom Erez, Serkan Cabi, Saran Tunyasuvunakool, János Kramár, Raia Hadsell, Nando de Freitas, et al. 2018. Reinforcement and imitation learning for diverse visuomotor skills. *arXiv preprint arXiv:1802.09564* (2018).
- [54] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. 2008. Maximum entropy inverse reinforcement learning. In *AAAI Conference on Artificial Intelligence*, Vol. 8. Chicago, IL, USA, 1433–1438.