Curriculum Reinforcement Learning using Optimal Transport via Gradual Domain Adaptation

Peide Huang, Mengdi Xu, Jiacheng Zhu, Laixi Shi, Fei Fang, Ding Zhao
Carnegie Mellon University
Pittsburgh, PA 15213
{peideh, mengdixu, jzhu4, laixis, feifang, dingzhao}@andrew.cmu.edu

Abstract

Curriculum Reinforcement Learning (CRL) aims to create a sequence of tasks, starting from easy ones and gradually learning towards difficult tasks. In this work, we focus on the idea of framing CRL as interpolations between a *source* (auxiliary) and a *target* task distribution. Although existing studies have shown the great potential of this idea, it remains unclear how to formally quantify and generate the movement between task distributions. Inspired by the insights from gradual domain adaptation in semi-supervised learning, we create a natural curriculum by breaking down the potentially large task distributional shift in CRL into smaller shifts. We propose GRADIENT, which formulates CRL as an optimal transport problem with a tailored distance metric between tasks. Specifically, we generate a sequence of task distributions as a geodesic interpolation (i.e., Wasserstein barycenter) between the source and target distributions. Different from many existing methods, our algorithm considers a task-dependent contextual distance metric and is capable of handling nonparametric distributions in both continuous and discrete context settings. In addition, we theoretically show that GRADIENT enables smooth transfer between subsequent stages in the curriculum under certain conditions. We conduct extensive experiments in locomotion and manipulation tasks and show that our proposed GRADIENT achieves higher performance than baselines in terms of learning efficiency and asymptotic performance.

1 Introduction

Reinforcement Learning (RL) [1] has demonstrated great potential in solving complex decision-making tasks [2], including but not limited to video games [3], chess [4], and robotic manipulation [5]. Among them, various prior works highlight daunting challenges resulting from sparse rewards. For example, in the maze navigation task, since the agent needs to navigate from the initial position to the goal to receive a positive reward, the task requires a large amount of randomized exploration. One solution to address this issue is Curriculum Reinforcement Learning (CRL) [6, 7], of which the objective is to create a sequence of environments to facilitate the learning of difficult tasks.

Although there are different interpretations of CRL, we focus on the one that views a curriculum as a sequence of task distributions that interpolate between a *source* (auxiliary) task distribution and a *target* task distribution [8, 9, 10, 11, 12, 13]. This interpretation allows more general presentations of tasks and a wider range of objectives such as generalization (uniform distribution over the task collection) or learning to complete difficult tasks (a subset of the task collection). Again we use the maze navigation as an example. Given a fixed maze layout and goal, a task is then defined by a start position, and the task distribution is a categorical distribution over all possible start positions. With a target distribution putting mass over start positions far away from the goal position, a natural

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

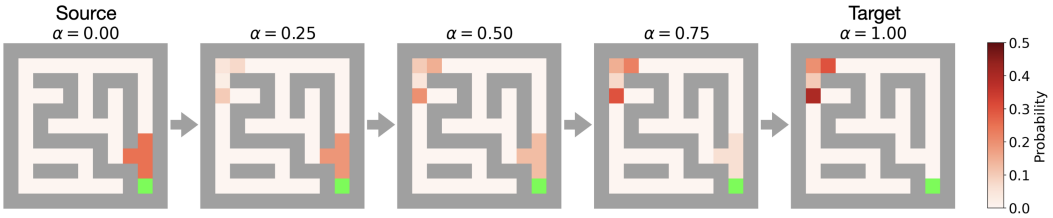

(a) Linear interpolation

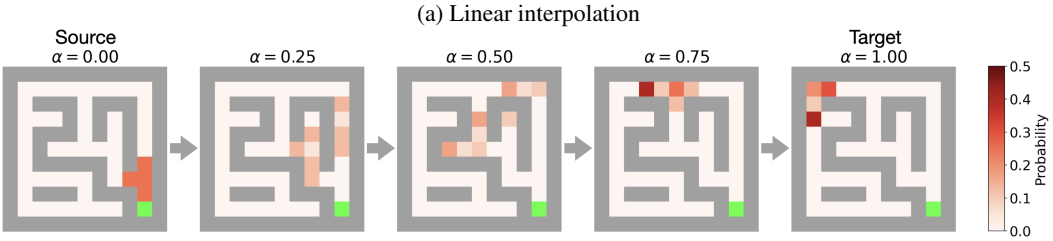

(b) GRADIENT interpolation (ours)

Figure 1: Intermediate task distributions generated by (a) linear interpolation and (b) our method for the maze navigation task. The green cell represents the goal. The red cells represent the initial positions (the darker the color is, the higher the probability is). The first column shows the source task distribution and the last column shows the target task distribution. In (a), linear interpolations do not cover cells where the source and the target have zero probability, which hardly benefits the learning. In contrast, (b) GRADIENT interpolations creates a curriculum that gradually morphs from the source to the target, covering tasks of intermediate difficulty and improving the learning efficiency.

curriculum to accelerate the learning process is putting start positions close to the goal position first, and gradually moving them towards the target distribution.

However, most of the existing methods, that interpret the curriculum as shifting distributions, use Kullback–Leibler (KL) divergence to measure the distance between distributions. This setting imposes several restrictions. First, due to either problem formulations or the computational feasibility, existing methods often require the distribution to be parameterized, e.g., Gaussian [8, 9, 10, 13], which limits the usage in practice. Second, most of the existing algorithms using KL divergence implicitly assume an l_2 Euclidean space which ignores the manifold structure when parameterizing RL environments [14].

In light of the aforementioned issues with the existing CRL method, we propose GRADIENT, an algorithm that creates a sequence of task distributions gradually morphing from the source to the target distribution using Optimal Transport (OT). GRADIENT approaches CRL from a gradual domain adaptation (GDA) perspective, breaking the potentially large domain shift between the source and the target into smaller shifts to enable efficient and smooth policy transfer. In this work, we first define a distance metric between individual tasks. Then we can find a series of task distributions that interpolate between the easy and the difficult task distribution by computing the Wasserstein barycenter. GRADIENT is able to deal with both discrete and continuous environment parameter spaces, and nonparametric distributions (represented either by explicit categorical distributions or implicit empirical distributions of particles). Under some conditions [15], GRADIENT provably ensures a smooth adaptation from one stage to the next.

We summarize our main contributions as follows:

1. We propose GRADIENT, a novel CRL framework based on optimal transport to generate gradually morphing intermediate task distributions. As a result, GRADIENT requires little effort to transfer between subsequent stages and therefore improves the learning efficiency towards difficult tasks.
2. We develop π-contextual-distance to measure the task similarity and compute the Wasserstein barycenters as intermediate task distributions. Our proposed method is able to deal with both continuous and discrete context spaces as well nonparametric distributions. We also prove the theoretical bound of policy transfer performance which leads to practical insights.
3. We demonstrate empirically that GRADIENT has stronger learning efficiency and asymptotic performance in a wide range of locomotion and manipulation tasks when compared with state-of-the-art CRL baselines.

2 Related Work

Curriculum reinforcement learning. Curriculum reinforcement learning (CRL) [6, 16] focuses on the generation of training environments for RL agents. There are several objectives in CRL: improving learning efficiency towards difficult tasks (*time-to-threshold*), maximum return (*asymptotic performance*), or transfer policies to solve unseen tasks (*generalization*). From a domain randomization perspective, Active Domain Randomization [5, 17] uses curricula to diversify the physical parameters of the simulator to facilitate the generalization in sim-to-real transfer. From a game-theoretical perspective, adversarial training is also developed to improve the robustness of RL agents in unseen environments [18, 19, 20, 21]. From an intrinsic motivation perspective, methods have been proposed to create curricula even in the absence of a target task to be accomplished [22, 13, 23].

CRL as an interpolation of distributions. In this work, we focus on another stream of works that interprets CRL as an explicit interpolation between an auxiliary task distribution and a difficult task distribution [8, 9, 10, 11]. Self-Paced Reinforcement Learning (SPRL) [8] is proposed to generate intermediate distributions by measuring the task distribution similarity using Kullback–Leibler (KL) divergence. However, as we will show in this paper, KL divergence brings several shortcomings that may impede the usage of the algorithms. First, although the formulation of [8, 9, 10, 11] does not restrict the distribution class, the algorithm realization requires the explicit computation of KL divergence, which is analytically tractable only for a restricted family of distributions. Second, using KL divergence implicitly assumes a l_2 Euclidean space which ignores the manifold structure when parameterizing RL environments. In this work, we use Wasserstein distance instead of KL divergence to measure the distance between distributions. Unlike KL divergence, Wasserstein distance considers the ground metric information and opens up a wide variety of task distance measures.

CRL using Optimal Transport. Hindsight Goal Generation (HGG) [24] aims to solve the poor exploration problem in Hindsight Experience Replay (HER). HGG computes 2-Wasserstein barycenter approximately to guide the hindsight goals towards the target distribution in an implicit curriculum. Concurrent to our work, CURROT [25] also uses optimal transport to generate intermediate tasks explicitly. CURROT formulates CRL as a constrained optimization problem with 2-Wasserstein distance to measure the distance between distributions. The main difference is that we propose task-dependent contextual distance metrics and directly treat the interpolation as the geodesic from the source to the target distribution.

Gradual domain adaptation in semi-supervised learning Gradual domain adaptation (GDA) [26, 27, 28, 29, 30, 31, 32, 33] considers the problem of transferring a classifier trained in a source domain with labeled data, to a target domain with unlabelled data. GDA solves this problem by designing a sequence of learning tasks. The classifier is retrained with the pseudolabels created by the classifier from the last stage in the sequence. Most of the existing literature assumes that there exist intermediate domains. However, there are a few works aiming to tackle the problem when intermediate domains, or the index (i.e., stage in the curriculum), are not readily available. A coarse-to-fine framework is proposed to sort and index intermediate domain data [33]. Another study proposes to create virtual samples from intermediate distributions by interpolating representations of examples from source and target domains and suggests using the optimal transport map to create interpolate data in semi-supervised learning [32]. It is demonstrated theoretically in [27] that the optimal path of samples is the geodesic interpolation defined by the optimal transport map. Our work is inspired by the *divide and conquer* paradigm in GDA and also uses the geodesic as our curriculum plan (although in a different learning paradigm).

3 Preliminary

3.1 Contextual Markov Decision Process

A contextual Markov decision process (CMDP) extends the standard single-task MDP to a multi-task setting. In this work, we consider discounted infinite-horizon CMDPs, represented by a tuple $M = (\mathcal{S}, \mathcal{C}, \mathcal{A}, R, P, p_0, \rho, \gamma)$. Here, $\mathcal{S}$ is the state space, $\mathcal{C}$ is the context space, $\mathcal{A}$ is the action space, $R : \mathcal{S} \times \mathcal{A} \times \mathcal{C} \mapsto \mathbb{R}$ is the context-dependent reward function, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{C} \mapsto \Delta(\mathcal{S})$ is the context-dependent transition function, $p_0 : \mathcal{C} \mapsto \Delta(\mathcal{S})$ is the context-dependent initial state distribution, $\rho \in \Delta(\mathcal{C})$ is the context distribution and $\gamma \in (0, 1)$ is the discount factor. Note that goal-conditioned reinforcement learning [12] can be considered as a special case of the CMDP.

To sample a trajectory $\boldsymbol{\tau} := \{s_t, a_t, r_t\}_{t=0}^{\infty}$ in CMDPs, the context $c \sim \rho$ is randomly generated by the environment at the beginning of each episode. With the initial state $s_0 \sim p_0(\cdot \mid c)$, at each time step t, the agent follows a policy π to select an action $a_t \sim \pi(s_t, c)$ and receives a reward $R(s_t, a_t, c)$. Then the environment transits to the next state $s_{t+1} \sim P(\cdot \mid s_t, a_t, c)$. Contextual reinforcement learning naturally extends the original RL objective to include the context distribution ρ. To find the optimal policy, we need to solve the following optimization problem:

$$\max_{\pi} V^{\pi}(\rho) = \max_{\pi} \mathbb{E}_{\boldsymbol{\tau}} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, c) \;\middle|\; c \sim \rho; \pi \right] \tag{1}$$

3.2 Optimal Transport

Wasserstein distance. The Kantorovich problem [34], a classic problem in optimal transport [35], aims to find the optimal coupling θ^* which minimizes the transportation cost between measures $\mu, \nu \in \mathcal{M}(\mathcal{C})$. Therefore, the *Wasserstein distance* defines the distance between probability distributions:

$$\mathcal{W}_d(\mu, \nu) = \inf_{\theta \in \Theta(\mu,\nu)} \int_{\mathcal{C}\times\mathcal{C}} d(c_s, c_t)\, \mathrm{d}\theta(c_s, c_t), \text{ subject to } \Theta = \left\{\theta : \gamma^{\mathcal{C}}_{\#}\theta = \mu, \gamma^{\mathcal{C}}_{\#}\theta = \nu\right\} \tag{2}$$

where $\mathcal{C}$ is the support space, $\Theta(\mu, \nu)$ is the set of all couplings between μ and ν, $d(\cdot, \cdot) : \mathcal{C}\times\mathcal{C} \mapsto \mathbb{R}_{\geq 0}$ is a distance function, $\gamma^{\mathcal{C}}$ is the projection from $\mathcal{C} \times \mathcal{C}$ onto $\mathcal{C}$, and $T_{\#}P$ generally denotes the push-forward measure of P by a map T [36, 37, 38]. This optimization is well-defined and the optimal θ always exists under mild conditions [35].

Wasserstein Geodesic and Wasserstein Barycenter. To construct a curriculum that allows the agent to efficiently solve the difficult task distribution, we follow the *Wasserstein geodesic* [39], the shortest path [27] under the Wasserstein distance between the source and the target distributions. While the original Wasserstein barycenter [40, 41] problem focuses on the Frèchet mean of multiple probability measures in a space endowed with the Wasserstein metric, we consider only two distributions μ_0 and μ_1. The set of barycenters between μ_0 and μ_1 is the geodesic curve given by McCann's interpolation [42, 43]. Thus, the interpolation between two given distributions μ_0 and μ_1 is defined as:

$$\nu_\alpha = \arg\min_{\nu'_\alpha} (1-\alpha)\mathcal{W}_d(\mu_0, \nu'_\alpha) + \alpha \mathcal{W}_d(\nu'_\alpha, \mu_1), \tag{3}$$

where each $\alpha \in [0, 1]$ specifies one unique interpolation distribution on the geodesic. While the computational cost of Wasserstein distance objective in Equation (2) could be a potential obstacle, we can follow the entropic optimal transport and utilize the celebrated Sinkhorn's algorithm [44]. Moreover, we can adopt a smoothing bias to solve for scalable debiased Sinkhorn barycenters [45].

4 Curriculum Reinforcement Learning using Optimal Transport

We first formulate the curriculum generation as an interpolation problem between probability distributions in Section 4.1. Then we propose a distance measure between contexts in Section 4.2 in order to compute the Wasserstein barycenter. Next, we show our main algorithm, GRADIENT, in Section 4.3 and an associated theorem that provides practical insights in Section 4.4.

4.1 Problem Formulation

Formally, given a target task distribution $\nu(c) \in \Delta(\mathcal{C})$, we aim to automatically generate a curriculum of task distributions, $\rho_0(c), \rho_1(c), \ldots, \rho_K(c)$ with K stages, that enables the agent to gradually adapt from the auxiliary task distribution $\mu(c)$ to the target $\nu(c)$, i.e., $\rho_0(c) \to \mu(c)$ and $\rho_K(c) \to \nu(c)$. If the context space $\mathcal{C}$ is discrete (sometimes called fixed-support in OT [35]) with cardinality $|\mathcal{C}|$, the task distribution is represented by a categorical distribution, e.g., $\mu(c) = [p(c_1), p(c_2), ..., p(c_{|\mathcal{C}|})]$, where $p(c_i) \geq 0, \sum_i p(c_i) = 1$. Whereas if the context space $\mathcal{C}$ is continuous (sometimes called free-support in OT), the task distribution is then approximated by a set of particles sampled from the distribution, e.g., $\mu(c) \approx \hat{\mu}(c) = \frac{1}{n_s}\sum_{i=1}^{n_s} \delta_{c_{si}}(c)$, where $\delta_{c_{si}}(c)$ is a Dirac delta at c_{si}. This highlights the capability of dealing with nonparametric distributions in our formulation and algorithms, in contrast to existing algorithms [8, 13, 8, 9, 10, 11].

4.2 Contextual Distance Metrics

To define the distance between contexts, we start by defining the distance between states.

Bisimulation metrics [46, 47, 48] measure states' *behaviorally equivalence* by defining a recursive relationship between states. Recent literature in RL [49, 14] uses the bisimulation concept to train state encoders that facilitate multi-tasking and generalization. [50] uses this bisimulation concept to enforce the policy to visit state-action pairs close to the support of logged transitions in offline RL.

However, the bisimulation metric is inherently "pessimistic" because it considers equivalence under all actions, even including the actions that never lead to positive outcomes for the agent [48]. To address this issue, [48] proposes on-policy π-bisimulation that considers the dynamics induced by the policy π. Similarly, we extend this notion to the CMDP settings and define the **contextual π-bisimulation metric**:

$$d^{\pi}_{c_i,c_j}(s_i, s_j) = |R^{\pi}_{s_i,c_i} - R^{\pi}_{s_j,c_j}| + \gamma \mathcal{W}_d(\mathcal{P}^{\pi}_{s_i,c_i}, \mathcal{P}^{\pi}_{s_j,c_j}) \tag{4}$$

where $R^{\pi}_{s,c} := \sum_a \pi(a \mid s) R(s, a, c)$ and $P^{\pi}_{s,c} := \sum_a P(\cdot \mid s, a, c)\pi(a \mid s)$. With the definition of contextual π-bisimulation metric, we are ready to propose the **π-contextual-distance** to measure the distance between two contexts:

Definition 4.1 (π-contextual-distance) *Given a CMDP $M = (\mathcal{S}, \mathcal{C}, \mathcal{A}, R, P, p_0, \rho, \gamma)$, the distance between two contexts $d^{\pi}(c_i, c_j)$ under the policy π is defined as*

$$d^{\pi}(c_i, c_j) = \mathbb{E}_{s_i \sim p_0(\cdot|c_i), s_j \sim p_0(\cdot|c_j)} \left[d^{\pi}_{c_i,c_j}(s_i, s_j) \right] \tag{5}$$

Conceptually, the **π-contextual-distance** approximately measures the performance difference between two contexts c_i and c_j under π. The algorithm to compute a simplified version of this metric under some conditions is detailed in the Appendix C.1. Note that it is difficult to compute this metric precisely in general. In practice, we can design and compute some surrogate contextual distance metrics, depending on the specific tasks. There are situations where it is reasonable to use l_2 as a surrogate metric when the contextual distance resembles the Euclidean space, such as in some goal-conditioned continuous environments [24]. In addition, although the contextual distance is not a strict metric, we can still use it as a ground metric in the OT computation. With the concepts of a contextual distance metric, we now introduce the algorithms to generate intermediate task distribution to enable the agent to gradually transfer from the source to the target task distribution.

4.3 Algorithms

We present our main algorithm in Algorithm 1. We introduce an interpolation factor $\Delta\alpha$ to decide the difference between two subsequent task distributions in the curriculum. Smaller $\Delta\alpha$ means a smaller difference between subsequent task distributions. In the algorithm, we treat $\Delta\alpha$ as a constant for simplicity but in effect, it can be scheduled or even adaptive, which we leave to future work. Note that the $\Delta\alpha$ is analogical to the KL divergence constraint in [9, 10, 8]. The main difference is that we use Wasserstein distance instead of KL divergence.

At the beginning of each stage in the curriculum, we add a $\Delta\alpha$ to the previous α (starting from 0). Then we pass the α into the function `ComputeBarycenter` (Algorithm 2) to generate the intermediate task distribution. Note that when $\alpha = 0$ and 1, the generated distribution are the source and target distribution respectively. In the `ComputeBarycenter` function, the computation method differs depending on whether the context is discrete or continuous. After generating the intermediate task distribution, we optimize the agent in this task distribution until the accumulative return G reaches the threshold $\bar{G}$. Then the curriculum enters the next stage and repeats the process until $\alpha = 1$. In other words, the path of the intermediate task distributions is the OT geodesic on the manifold defined by Wasserstein distance.

4.4 Theoretical Analysis

The proposed algorithm benefits from breaking the difficulty of learning in the target domain into multiple small challenges by designing a sequence of K stages that gradually morph towards the target. This motivates us to theoretically answer the following question:

Can we achieve smooth transfer to a new stage based on the optimal policy of the previous stage?

Algorithm 1: GRAdual **D**omain adaptation for curriculum re**I**nforc**E**ment lear**N**ing via optimal **T**ransport (**GRADIENT**)

Input: Source task distribution $\mu(c)$, target task distribution $\nu(c)$, interpolation factor $\Delta\alpha$, distance metric d, reward threshold $\bar{G}$, maximum number of stages K.

Initialize the agent policy π
for k *in* $0, 1, 2, ..., K$ **do**
  $\alpha \leftarrow \min(k * \Delta\alpha, 1)$;
  $\rho(c) \leftarrow$ `ComputeBarycenter`(μ, ν, α, d);
  do
    $G \leftarrow$ Optimize π in the task distribution $\rho(c)$ and return training reward;
  until $G > \bar{G}$;
Output: Agent policy π

Algorithm 2: ComputeBarycenter

Input: Source and target μ, ν, interpolation constant α, distance metric d.
$w_1 \leftarrow \alpha, w_2 \leftarrow 1-\alpha, (b_1, b_2), l \leftarrow \mathbf{1}$;
if $\mathcal{C}$ *is discrete* **then**
  $\lambda_1 \leftarrow \mu(c), \lambda_2 \leftarrow \nu(c)$
  Cost matrix:
  $\mathbf{C} := d(\cdot,\cdot) : \mathcal{C} \times \mathcal{C} \mapsto \mathbb{R}^{n_s \times n_t}$;
else
  $\lambda_1 = \frac{1}{n_s}\sum_i^{n_s} c_{si}, \lambda_2 = \frac{1}{n_t}\sum_j^{n_t} c_{tj}$;
  Cost matrix $\mathbf{C} : C_{ij} = d(c_{si}, c_{tj})$;
$\mathbf{K} = exp\left(-\mathbf{C}/\epsilon\right)$ `// Sinkhorn param` ϵ
while *not converge* **do**
  $a_1 \leftarrow (\lambda_1/\mathbf{K}b_1), a_2 \leftarrow (\lambda_2/\mathbf{K}b_2)$;
  $\rho \leftarrow l \odot \prod_{m=1}^{2}(\mathbf{K}^\top a_m)^{w_m}$;
  $b_1 \leftarrow \left(\rho/\mathbf{K}^\top a_1\right), b_2 \leftarrow \left(\rho/\mathbf{K}^\top a_2\right)$;
  $l \leftarrow \sqrt{l \odot (\rho/\mathbf{K}l)}$;
Output: ρ

Towards this, motivated by the maze navigation example shown in Figure 1, throughout this section, we focus on the collection of CMDPs which obeys the following assumption:

Assumption 4.1 (Consistency and homogeneity) *The initial state distribution is consistent with the associated context distribution, namely* $p_0(s \,|\, c) = \mathbb{1}(s = c)$. *In addition, the transition kernels and rewards of the CMDP are homogeneous with respect to the context. Specifically,*

$$P(s' \,|\, s, a, c_i) = P(s' \,|\, s, a, c_j), \quad R(s, a, c_i) = R(s, a, c_j), \qquad \forall c_i, c_j \in \mathcal{C}, s \in \mathcal{S}, a \in \mathcal{A}.$$

To continue, for k-th stage associated with the context distribution $\rho_k \in \Delta(\mathcal{C})$, we denote the optimal policy in k-th stage as $\pi_k^\star := \arg\max_\pi V^\pi(\rho_k)$. Armed with the above notations and assumptions, we are positioned to introduce our main theorem.

Theorem 4.1 *Let m be a problem-dependent positive constant. Consider any MDP under the assumption 4.1 and the situations further specified in Appendix A.1. For any two subsequent stages $k, k+1 \in [K]$, the optimal policies $\pi_k^\star, \pi_{k+1}^\star$ obey*

$$V^{\pi_{k+1}^\star}(\rho_{k+1}) - V^{\pi_k^\star}(\rho_{k+1}) \le m\mathcal{W}_{d^\star}(\rho_k, \rho_{k+1}). \tag{6}$$

Here, $d^\star$ is the distance metric defined on the metric space $(\mathcal{C}, d^\star)$ which obeys ($\mathcal{U}(\mathcal{C})$ denote the uniform distribution over $\mathcal{C}$)

$$\pi^\star := \max_\pi V^\pi\left(\mathcal{U}(\mathcal{C})\right), \qquad d^\star(\cdot,\cdot) := d^{\pi^\star}(\cdot,\cdot). \tag{7}$$

The proof can be found in Appendix A.4. In addition, the following corollary can be directly verified by Theorem 4.1.

Corollary 4.1.1 *By constant speed geodesics (7.2 of [51]), the stages generated by GRADIENT obey*

$$V^{\pi_{k+1}^\star}(\rho_{k+1}) - V^{\pi_k^\star}(\rho_{k+1}) \le m\Delta\alpha\mathcal{W}_{d^\star}(\mu, \nu). \tag{8}$$

Remark 4.1 *Theorem 4.1 combined with Corollary 4.1.1 indicates that with tailored stages by GRADIENT (sufficiently small $\Delta\alpha$), it is easy to adapt the learned optimal policy $\pi_k^\star$ in the current stage to the optimal policy $\pi_{k+1}^\star$ in the next stage defined by ρ_{k+1}, since the performance gap is small. In particular, after we obtain the optimal policy $\pi_k^\star$ associated with the k-th stage in the curriculum and transfer this policy to the next stage, the performance gap between it and the optimal policy in the $(k+1)$-th stage can be controlled by the Wasserstein distance between two subsequent context distributions and therefore a fraction of the Wasserstein distance between the source and the target (which can be seen as a constant). By properly controlling the Wasserstein distance using GRADIENT, we can achieve gradual domain adaptation from the source to the target task distribution.*

5 Experiments

Evaluation Metric. To quantify the benefit of CRL, we compare the learning progress of an agent evaluated on the target task distribution [1]. More concretely, the metric we use is *time-to-threshold* [6], which shows how fast an agent can learn a policy that achieves a satisfying return on the target task if it transfers knowledge. We also compare the *asymptotic performance* of the agent. For the learner, we use the SAC [52] and PPO [53] implementations provided in the `StableBaseline3` library [54].

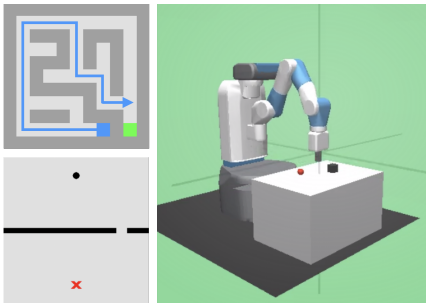

Figure 2: Environment visualizations. For the **Maze** (upper-left), the context is the initial state. For the **PointMass** (lower-left), the context is the position and width of the gap on the wall. For the **FetchPush**, the context is the goal position on a 2D surface.

Environments. The first **Maze** environment is to investigate the performance of `GRADIENT` and the effect of $\Delta\alpha$. Maze has a discrete context space and therefore the exact context distance metric is available. The next two environments are to investigate the benefits of `GRADIENT` when the context space is continuous and we use l_2 distance as a surrogate contextual distance. In the **PointMass** environment, the source and the target distribution are both Gaussian distributions, which are in line with the settings of existing algorithms. In the **FetchPush** environment, we want to investigate the case where the source and the target are not Gaussian distributions.

We benchmark our algorithm with six baselines including four CRL algorithms:

1. `No-Curriculum`: sampling contexts according to the target task distribution.
2. `Domain-Randomization`: uniformly sampling the contexts from the context space.
3. `Linear-Interpolation`: sampling contexts based on a weighted average between two distributions, i.e., $(1-\alpha)\mu + \alpha\nu$.
4. `ALP-GMM` [13]: the state-of-the-art intrinsically motivated CRL algorithm that maximizes the absolute learning progress of the agent.
5. `Goal-GAN` [12]: using GAN to generate contexts of intermediate difficulties and automatically explore the goal space.
6. `SPDL` [8] and `WB-SPDL` [55]: interpreting the curriculum generation as a variational inference problem to progressively approach the target task distribution. `WB-SPDL` extends SPDL by using Wasserstein distance instead of KL divergence.

5.1 Can GRADIENT Handle Discrete Contexts and What is the Effect of $\Delta\alpha$?

In the **Maze** task, we fix the layout of the maze and the goal. The state and action spaces are discrete. The action space includes movements towards four directions. The context is the initial position, and therefore the context space overlaps with the state space. The observation is the flattened value representation of the maze, including the goal, the current position, and the layout.

We visualize the optimal π^*-contextual-distance metric in Figure 3a (computed using the optimal policy and normalized to make the maximum value be 1). The distance metric can be represented by a symmetrical matrix. We then generate curricula using `GRADIENT` with $\Delta\alpha = 0.2, 0.1, 0.05$. From Figure 3b, our proposed method outperforms the baselines significantly in terms of the time-to-threshold evaluated in the target distribution. With a smaller $\Delta\alpha$, `GRADIENT` learns slightly slower; nevertheless, all choices achieve good performances. In this simple environment, `Domain-Randomization` also improves the learning efficiency compared with `No-Curriculum` (which cannot learn at all). Interestingly, `Domain-Randomization` even outperforms the `Linear-Interpolation`, which demonstrates the caveat of bad intermediate task distribution.

5.2 How Does GRADIENT Perform When the Distributions Are Gaussian?

In PointMass [8], the agent needs to navigate a pointmass through a wall with a small gap at an off-center position to reach the goal. The context is a 2-dimension vector representing the width and

[1] Code is available under https://github.com/PeideHuang/gradient.git

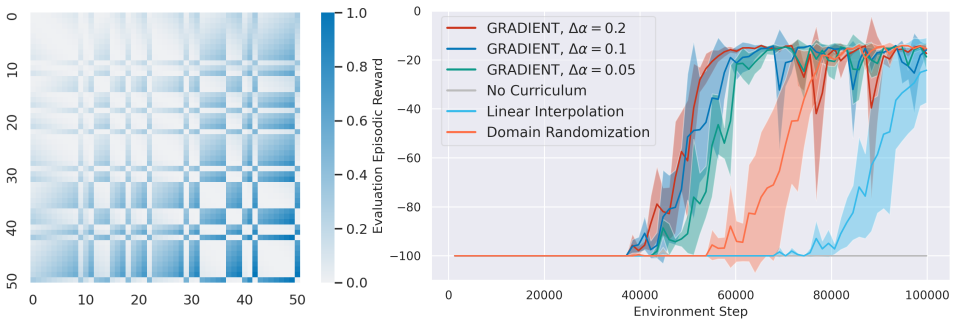

(a) π-contextual-distance (b) Evaluation reward in the target

Figure 3: The π-contextual-distance and evaluation results in Maze. The shaded area represents the standard deviation. We use a PPO learner and show that GRADIENT converges faster than baselines.

position of the gap on the wall. Following the setting in the SPDL paper [8], the target distribution is an isotropic Gaussian distribution centered at $[2.5, 0.5]$ with a negligible variance $[4\times10^{-3}, 3.75\times10^{-3}]$ (which is effectively a point as shown in Figure 5a). The source distribution is an isotropic Gaussian distribution centered at $[0, 4.25]$ with a variance of $[2, 1.875]$.

From Figure 4a, we observe that both GRADIENT and SPDL significantly outperform other baselines in terms of the time-to-threshold and achieve the same asymptotic performance. It is reasonable since these two methods are capable of specifying the target distribution to guide curriculum generation. We present the curricula visualization of both methods in Figure 5a and show that the intermediate task distributions gradually move from the source to the target distribution. However, the other three baselines lack the flexibility to leverage the target distribution. This experiment highlights the importance of being able to specify the target distribution.

5.3 How Does GRADIENT Perform When the Distributions Are NOT Gaussian?

In FetchPush [56], the objective is to use the gripper to push the box to a goal position. The observation space is a 28-dimension vector, including information about the goal. The context is a 2-dimension vector representing the goal position on a surface. The target distribution is a uniform distribution over the circumference of a half-circle (Figure 5b). The source distribution is a uniform distribution over a square region centered at the box position, excluding the region within a certain radius of the object. We use this experiment to highlight the importance of the capability to handle arbitrary distributions rather than only the parametric Gaussian distributions. Since SPDL can only deal with parametric Gaussian distribution, we first fit the target and the source distribution with two Gaussian distributions and feed them into the baseline algorithms.

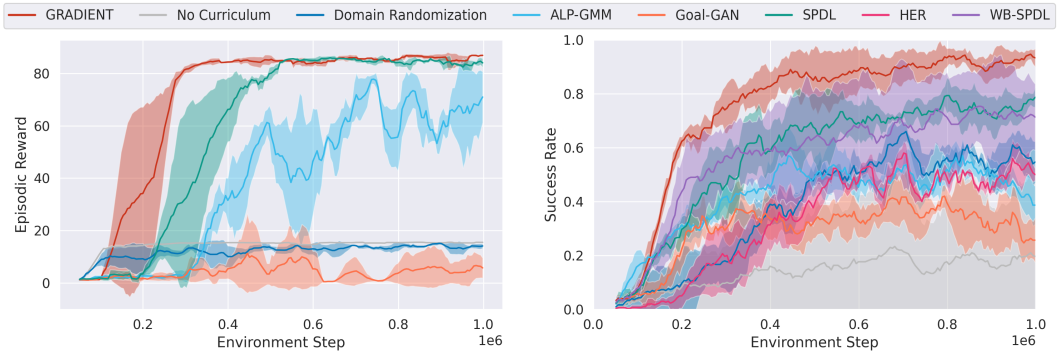

(a) PointMass (b) FetchPush

Figure 4: Evaluation performance in the target task distribution (with a rolling average of 10) for PointMass and FetchPush. In PointMass, where both the source and target distributions are defined by a parametric Gaussian, GRADIENT converges faster than baselines. In FetchPush where the target task distribution can NOT be accurately approximated by a Gaussian, GRADIENT converges faster and has a better asymptotic performance. We use SAC as the learner for both environments.

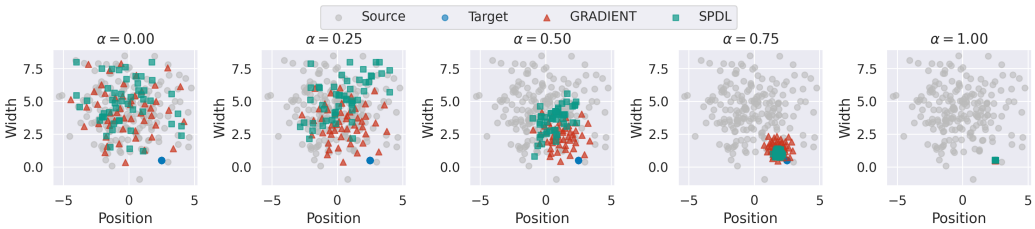

(a) PointMass context visualization

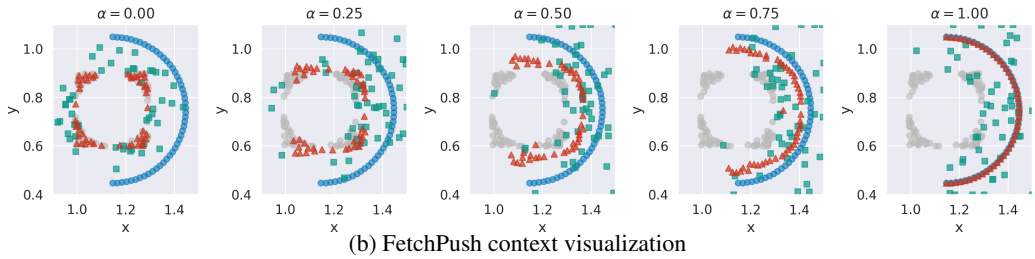

(b) FetchPush context visualization

Figure 5: Visualizations of context distributions and curricula in PointMass and FetchPush. We present GRADIENT's curricula corresponding to $\alpha = 0.00, 0.25, 0.50, 0.75, 1.00$ together with SPDL's contexts from 5 representative stages. For PointMass, the contexts of SPDL are taken from environment steps at 10k, 50k, 100k, 200k, and 300k. For FetchPush, the contexts of SPDL are taken from environment steps at 5k, 50k, 100k, 250k, and 500k.

From Figure 4b, GRADIENT outperforms all the baseline methods in terms of the time-to-threshold in the target distribution. For ALP-GMM and Goal-GAN, since these two methods do not consider the target distribution, they achieve lower performance than the target-aware methods such as SPDL and GRADIENT. We further compare these two target-aware methods and find that GRADIENT outperforms the SPDL in terms of learning efficiency and asymptotic performance. The reason can be explained by visualizing the context distribution in Figure 5b. SPDL indeed shows the behavior of moving the context distribution towards the target. However, since the target distribution cannot be fit with a Gaussian well, many sampled contexts are in fact far away from the target distribution, which may not improve the learning efficiency in the target domain. Although Domain-Randomization learns slightly lower than the CRL methods, it achieves even higher asymptotic performance than ALP-GMM and Goal-GAN. The potential reason could be that it does not bias towards a certain area while it is possible for ALP-GMM and Goal-GAN.

5.4 Beyond Euclidean Distance Metric: Using Distance Embedding for Continuous Spaces

In practice, it is not always appropriate to use the Euclidean distance as the surrogate contextual distance. Though there are many existing OT algorithms based on Euclidean distance metric, solving for free-support Wasserstein barycenters on non-Euclidean space is non-trivial [57, 58]. One solution to bypass this issue is to embed contexts to a latent space such that the Euclidean interpolation in the latent space approximates the interpolation in the original non-Euclidean space after reconstruction.

For the fixed-support setting, Deep Wasserstein Embedding (DWE) proposed in [59] uses Siamese networks [60] to learn a mapping function such that the Euclidean distance in the latent space approximates the Wasserstein distance in the original space. Motivated by this idea, we instead develop an embedding mechanism for free supports, i.e., continuous contexts. We leverage an encoder to embed contexts to latent space in which the Euclidean distance mimics the original distance metric. We train the encoder and decoder based on a reconstruction loss and a distance embedding loss. We include the details of distance embedding in Algorithm 4 and a modified version of GRADIENT in Algorithm 3 in Appendix C.2.

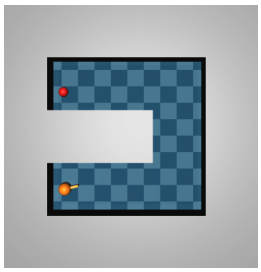

Figure 6: U-Maze environment. The orange ball represents the agent. The red ball represents the goal. The context is the goal position.

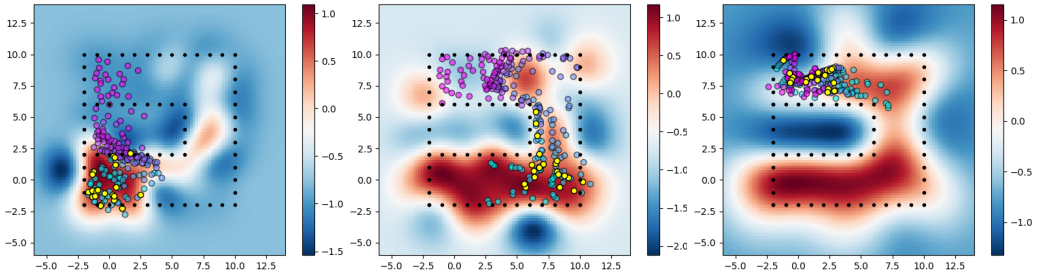

(a) Stage 1 (b) Stage 5 (c) Stage 9

Figure 7: U-Maze interpolation results. The original source and target distributions are two Gaussian distributions centered at $[0, 0]$ and $[0, 8]$. The color of the heat map represents the estimate of $J^\pi(c)$. The colored circles represent the Wasserstein interpolations. More specifically, the circles with cyan to purple color represent the potential interpolation results from the current source to target with $\alpha = 0.0, 0.1, 0.2, \ldots, 1.0$. The yellow circles highlight the selected barycenter corresponding to the current stage.

We show an example of using Algorithm 3. Here we consider a classical U-shaped maze with continuous spaces in Figure 6. We assume that the source and target distributions are two Gaussian distributions at the two ends of the maze. Due to the existence of the obstacle in the middle, it is not appropriate to use the l_2 distance as the contextual distance. In this case, we use $d^\pi(c_1, c_2) := |J^\pi(c_1) - J^\pi(c_2)|$ as a surrogate metric for the π-contextual-distance, where $J^\pi(c)$ represents the expected episodic reward of context c under the policy π.

The interpolation results are shown in Figure 7. At the first few stages, since the agent does not have a good policy or a good estimate of the episodic reward, the potential geodesic interpolation tries to go through the obstacle in the middle. Fortunately, since α is small at first, the selected intermediate barycenters (yellow circles) are relatively close to the source distribution, so `GRADIENT` will not generate unreasonable tasks. As the learning progressing, the estimates become better and better and therefore produce better interpolation results.

6 Conclusion and Limitation

In this work, we propose a novel curriculum reinforcement learning framework as an optimal transport problem. We also develop a distance metric to measure the difference between tasks. We show that our proposed method is able to achieve smooth transfer between subsequent stages in the generated curriculum. Finally, the empirical results highlight the three features of our methods: considering metric information, being aware of the target, and being able to handle nonparametric distributions. For future work, it could be interesting to investigate how to use the *intrinsic curiosity* driven type of method as in [22, 13, 23] to automatically generate the source distribution. In addition, wider choices of the context distance can be explored to provide more insights into the OT-based CRL methods. From a theoretical perspective, OT also provides powerful mathematical tools for future researches to investigate rigorously why CRL methods can help reduce the sample complexity.

There are two main limitations of our proposed GRADIENT. First, due to the limitation of OT computation complexity, we only examine the settings that are either discrete or low-dimensional. Such experiment design choice is common in the related works and baselines [13, 12, 8]. It comes from both the fact of many tasks can be represented by a low-dimensional vector, such as goal positioned or physical parameters of robots. We believe how to deal with high-dimensional contexts, such as images, is a promising future direction. Second, our analysis is conducted based on a restrictive subset of contextual MDPs (Assumption 4.1)where the context controls only the initial state distribution. It would be beneficial to relax this assumption by imposing assumptions on the specific manner in which the context controls the dynamics or reward.

Acknowledgments

We gratefully acknowledge support from the National Science Foundation under grant CAREER CNS-2047454. Fei Fang was supported in part by NSF grant IIS-2046640 (CAREER).

References

[1] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[2] Mengdi Xu, Wenhao Ding, Jiacheng Zhu, Zuxin Liu, Baiming Chen, and Ding Zhao. Task-agnostic online reinforcement learning with an infinite mixture of gaussian processes. *Advances in Neural Information Processing Systems*, 33:6429–6440, 2020.

[3] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[4] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.

[5] 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.

[6] Sanmit Narvekar, Bei Peng, Matteo Leonetti, Jivko Sinapov, Matthew E Taylor, and Peter Stone. Curriculum learning for reinforcement learning domains: A framework and survey. *Journal of machine learning research: JMLR*, 21(181):1–50, 2020.

[7] Mengdi Xu, Zuxin Liu, Peide Huang, Wenhao Ding, Zhepeng Cen, Bo Li, and Ding Zhao. Trustworthy reinforcement learning against intrinsic vulnerabilities: Robustness, safety, and generalizability. *arXiv preprint arXiv:2209.08025*, 2022.

[8] Pascal Klink, Hany Abdulsamad, Boris Belousov, Carlo D'Eramo, Jan Peters, and Joni Pajarinen. A probabilistic interpretation of self-paced learning with applications to reinforcement learning. *Journal of Machine Learning Research*, 22(182):1–52, 2021.

[9] Pascal Klink, Carlo D'Eramo, Jan R Peters, and Joni Pajarinen. Self-paced deep reinforcement learning. *Advances in Neural Information Processing Systems*, 33:9216–9227, 2020.

[10] Pascal Klink, Hany Abdulsamad, Boris Belousov, and Jan Peters. Self-paced contextual reinforcement learning. In *Conference on Robot Learning*, pages 513–529. PMLR, 2020.

[11] Jiayu Chen, Yuanxin Zhang, Yuanfan Xu, Huimin Ma, Huazhong Yang, Jiaming Song, Yu Wang, and Yi Wu. Variational automatic curriculum learning for sparse-reward cooperative multi-agent problems. *Advances in Neural Information Processing Systems*, 34, 2021.

[12] Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. In *International conference on machine learning*, pages 1515–1528. PMLR, 2018.

[13] 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.

[14] Philippe Hansen-Estruch, Amy Zhang, Ashvin Nair, Patrick Yin, and Sergey Levine. Bisimulation makes analogies in goal-conditioned reinforcement learning. *arXiv preprint arXiv:2204.13060*, 2022.

[15] Gen Li, Laixi Shi, Yuxin Chen, Yuantao Gu, and Yuejie Chi. Breaking the sample complexity barrier to regret-optimal model-free reinforcement learning. *Advances in Neural Information Processing Systems*, 34:17762–17776, 2021.

[16] Clément Romac, Rémy Portelas, Katja Hofmann, and Pierre-Yves Oudeyer. Teachmyagent: a benchmark for automatic curriculum learning in deep rl. In *International Conference on Machine Learning*, pages 9052–9063. PMLR, 2021.

[17] 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.

[18] Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, pages 2817–2826. PMLR, 2017.

[19] Michael Dennis, Natasha Jaques, Eugene Vinitsky, 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.

[20] Peide Huang, Mengdi Xu, Fei Fang, and Ding Zhao. Robust reinforcement learning as a stackelberg game via adaptively-regularized adversarial training. *arXiv preprint arXiv:2202.09514*, 2022.

[21] Mengdi Xu, Peide Huang, Fengpei Li, Jiacheng Zhu, Xuewei Qi, Kentaro Oguchi, Zhiyuan Huang, Henry Lam, and Ding Zhao. Accelerated policy evaluation: Learning adversarial environments with adaptive importance sampling. *arXiv preprint arXiv:2106.10566*, 2021.

[22] Adrien Baranes and Pierre-Yves Oudeyer. Intrinsically motivated goal exploration for active motor learning in robots: A case study. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1766–1773. IEEE, 2010.

[23] Pierre Fournier, Olivier Sigaud, Mohamed Chetouani, and Pierre-Yves Oudeyer. Accuracy-based curriculum learning in deep reinforcement learning. *arXiv preprint arXiv:1806.09614*, 2018.

[24] Zhizhou Ren, Kefan Dong, Yuan Zhou, Qiang Liu, and Jian Peng. Exploration via hindsight goal generation. *Advances in Neural Information Processing Systems*, 32, 2019.

[25] Pascal Klink, Haoyi Yang, Carlo D'Eramo, Jan Peters, and Joni Pajarinen. Curriculum reinforcement learning via constrained optimal transport. In *International Conference on Machine Learning*, pages 11341–11358. PMLR, 2022.

[26] Ananya Kumar, Tengyu Ma, and Percy Liang. Understanding self-training for gradual domain adaptation. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5468–5479. PMLR, 13–18 Jul 2020.

[27] Haoxiang Wang, Bo Li, and Han Zhao. Understanding gradual domain adaptation: Improved analysis, optimal path and beyond. *arXiv preprint arXiv:2204.08200*, 2022.

[28] Judy Hoffman, Trevor Darrell, and Kate Saenko. Continuous manifold based adaptation for evolving visual domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 867–874, 2014.

[29] Michael Gadermayr, Dennis Eschweiler, Barbara Mara Klinkhammer, Peter Boor, and Dorit Merhof. Gradual domain adaptation for segmenting whole slide images showing pathological variability. In *International Conference on Image and Signal Processing*, pages 461–469. Springer, 2018.

[30] Markus Wulfmeier, Alex Bewley, and Ingmar Posner. Incremental adversarial domain adaptation for continually changing environments. In *2018 IEEE International conference on robotics and automation (ICRA)*, pages 4489–4495. IEEE, 2018.

[31] Hao Wang, Hao He, and Dina Katabi. Continuously indexed domain adaptation. *arXiv preprint arXiv:2007.01807*, 2020.

[32] Samira Abnar, Rianne van den Berg, Golnaz Ghiasi, Mostafa Dehghani, Nal Kalchbrenner, and Hanie Sedghi. Gradual domain adaptation in the wild: When intermediate distributions are absent. *arXiv preprint arXiv:2106.06080*, 2021.

[33] Hong-You Chen and Wei-Lun Chao. Gradual domain adaptation without indexed intermediate domains. *Advances in Neural Information Processing Systems*, 34, 2021.

[34] Leonid V Kantorovich. On the translocation of masses. *Journal of mathematical sciences*, 133(4):1381–1382, 2006.

[35] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.

[36] Jiacheng Zhu, Aritra Guha, Dat Do, Mengdi Xu, XuanLong Nguyen, and Ding Zhao. Functional optimal transport: map estimation and domain adaptation for functional data. *arXiv preprint arXiv:2102.03895*, 2021.

[37] Yiwen Dong, Jiacheng Zhu, and Hao Young Noh. Re-vibe: Vibration-based indoor person re-identification through cross-structure optimal transport. *1st International Workshop on The Future of Work, Workplaces, and Smart Buildings (FoWSB)*, 2022.

[38] Jiacheng Zhu, Gregory Darnell, Agni Kumar, Ding Zhao, Bo Li, Xuanlong Nguyen, and Shirley You Ren. Physiomtl: Personalizing physiological patterns using optimal transport multi-task regression. In *Conference on Health, Inference, and Learning*, pages 354–374. PMLR, 2022.

[39] Vivien Seguy and Marco Cuturi. Principal geodesic analysis for probability measures under the optimal transport metric. *Advances in Neural Information Processing Systems*, 28, 2015.

[40] Martial Agueh and Guillaume Carlier. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.

[41] Jiaojiao Fan, Amirhossein Taghvaei, and Yongxin Chen. Scalable computations of wasserstein barycenter via input convex neural networks. *arXiv preprint arXiv:2007.04462*, 2020.

[42] Robert J McCann. A convexity principle for interacting gases. *Advances in mathematics*, 128(1):153–179, 1997.

[43] Jiacheng Zhu, Jielin Qiu, Zhuolin Yang, Douglas Weber, Michael A Rosenberg, Emerson Liu, Bo Li, and Ding Zhao. Geoecg: Data augmentation via wasserstein geodesic perturbation for robust electrocardiogram prediction. *arXiv preprint arXiv:2208.01220*, 2022.

[44] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.

[45] Hicham Janati, Marco Cuturi, and Alexandre Gramfort. Debiased sinkhorn barycenters. In *International Conference on Machine Learning*, pages 4692–4701. PMLR, 2020.

[46] Norm Ferns, Prakash Panangaden, and Doina Precup. Bisimulation metrics for continuous markov decision processes. *SIAM Journal on Computing*, 40(6):1662–1714, 2011.

[47] Norman Ferns and Doina Precup. Bisimulation metrics are optimal value functions. In *UAI*, pages 210–219. Citeseer, 2014.

[48] Pablo Samuel Castro. Scalable methods for computing state similarity in deterministic markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10069–10076, 2020.

[49] Amy Zhang, Rowan McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning invariant representations for reinforcement learning without reconstruction. *arXiv preprint arXiv:2006.10742*, 2020.

[50] Robert Dadashi, Shideh Rezaeifar, Nino Vieillard, Léonard Hussenot, Olivier Pietquin, and Matthieu Geist. Offline reinforcement learning with pseudometric learning. In *International Conference on Machine Learning*, pages 2307–2318. PMLR, 2021.

[51] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005.

[52] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.

[53] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[54] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021.

[55] Pascal Klink, Haoyi Yang, Jan Peters, and Joni Pajarinen. Metrics matter: A closer look on self-paced reinforcement learning. 2021.

[56] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

[57] Marco Cuturi and Arnaud Doucet. Fast computation of wasserstein barycenters. In *International conference on machine learning*, pages 685–693. PMLR, 2014.

[58] Lingxiao Li, Aude Genevay, Mikhail Yurochkin, and Justin M Solomon. Continuous regularized wasserstein barycenters. *Advances in Neural Information Processing Systems*, 33:17755–17765, 2020.

[59] Nicolas Courty, Rémi Flamary, and Mélanie Ducoffe. Learning wasserstein embeddings. *arXiv preprint arXiv:1710.07457*, 2017.

[60] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 539–546. IEEE, 2005.

[61] Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite markov decision processes. In *UAI*, volume 4, pages 162–169, 2004.

[62] Yuxin Deng and Wenjie Du. The kantorovich metric in computer science: A brief survey. *Electronic Notes in Theoretical Computer Science*, 253(3):73–82, 2009.

[63] Xinfeng Zhou. *A Practical Guide to Quantitative Finance Interviews*. Lulu. com, 2008.

A Proof of Theorem 4.1

In this section, we shall provide the proof for Theorem 4.1. The rest of this section is organized as follows. Appendix A.1 provides additional useful notations and definitions including but not limited to CMDPs, value function, and distance metrics. Appendix A.2 introduces further assumptions, and A.3 introduces the preliminary of optimal transport. Then, we provide the proof pipeline of Theorem 4.1 in Appendix A.4 and postpone the auxiliary proof to Appendix A.5. Finally, we summarize all the useful notations in Table 1.

A.1 Additional notation and definitions used in the proof

Before starting, let's introduce some additional notations useful throughout the theoretical analysis. For any discrete set X, we will denote the set of probability measures on X by $\Delta(X)$. For any vector $x \in \mathbb{R}^{SA}$ (resp. $x \in \mathbb{R}^{S}$ or $x \in \mathbb{R}^{C}$) that constitutes certain values for each state-action pair (resp. state or context), we use $x(s,a)$ (resp. $x(s)$ or $x(c)$) to denote the entry associated with the (s,a) pair (resp state s or context c). We denote $\mathsf{supp}(\rho)$ as the support of any distribution ρ.

Target CMDPs. Throughout the proof, we shall focus on the set of CMDPs introduced in Assumption 4.1. Moreover, without loss of generality, let the state space $\mathcal{S} = \{1, 2, \cdots, S\}$ (resp. action space $\mathcal{A} = \{1, 2, \cdots, A\}$) to be with finite cardinality S (resp. A). Without loss of generality, let the immediate reward $R(s,a,s) \in [0, 1-\gamma]$ for all $(s,a) \in \mathcal{S} \times \mathcal{A}$.

First, we introduce the following *occupancy distributions* associated with policy π and any CMDP with the initial state distribution $c \sim \rho$:

$$o^{\pi}(s;\rho) \coloneqq (1-\gamma)\sum_{t=0}^{\infty}\gamma^{t}\mathbb{P}\big(s_t = s \,|\, c \sim \rho; \pi\big), \tag{9a}$$

$$o^{\pi}(s,a;\rho) \coloneqq (1-\gamma)\sum_{t=0}^{\infty}\gamma^{t}\mathbb{P}\big(s_t = s, a_t = a \,|\, c \sim \rho; \pi\big) = o^{\pi}(s;\rho)\pi(a \,|\, s). \tag{9b}$$

With this in mind, we introduce two important definitions in the proof,: the covering state-action (resp. state) space by executing π with the initial state distribution $c \sim \rho$ defined as

$$\mathcal{C}(\rho,\pi) = \{(s,a) : o^{\pi}(s,a;\rho) > 0\}, \qquad \mathcal{C}_s(\rho,\pi) = \{s : o^{\pi}(s;\rho) > 0\}. \tag{10}$$

The covering space represent the area that is possible to be visited by the policy π given the initial state distribution.

Value function. Different from the value function defined by context distribution ρ, with abuse of notation, we denote the value function V_{s} associated with a state s, a policy π and the target CMDPs as

$$V_{\mathsf{s}}^{\pi}(s) = \mathbb{E}\left[\sum_{t=0}^{\infty}\gamma^{t}R(s_t,a_t,c') \;\middle|\; s_0 = s, c = s; \pi\right]. \tag{11}$$

For convenience, for any $s \in \mathcal{S}$, we define $\rho_s(c) \coloneqq \mathbb{1}(c = s)$. Recalling that the target MDPs are defined by $p_0(s \,|\, c) = \mathbb{1}(s = c)$, it is easily verified that for any policy π, the value function satisfies the following properties:

$$\forall s \in \mathcal{S}: \qquad 0 \le V_{\mathsf{s}}^{\pi}(s) = V^{\pi}(\rho_s) \le 1, \tag{12a}$$

$$\forall \rho \in \Delta(\mathcal{C}) \in \Delta(\mathcal{S}): \qquad V_{\mathsf{s}}^{\pi}(\rho) = V^{\pi}(\rho). \tag{12b}$$

Distance metrics. Furthermore, for the specified target CMDPs, plugging in $p_0(s \,|\, c) = \mathbb{1}(s = c)$, for any two contexts $c_i, c_j \in \mathcal{C}$, the **π-contextual-distance** metric defined in context space $\mathcal{C}$ obeys

$$d^{\pi}(c_i,c_j) = \mathbb{E}_{s_i \sim p_0(\cdot|c_i), s_j \sim p_0(\cdot|c_j)}\left[d^{\pi}_{c_i,c_j}(s_i,s_j)\right] = d^{\pi}_{s_i,s_j}(s_i,s_j), \tag{13}$$

where s_i (resp. s_j) is the corresponding initial state when the context are c_i (resp. c_j). With this observation, we highlight that for the target CMDPs, the **contextual π-bisimulation** metric degrades to **π-bisimulation** $d^{\pi}_{\mathsf{s}}(\cdot,\cdot)$ [48] such that:

$$\forall s_i, s_j \in \mathcal{S}: \qquad d^{\pi}_{\mathsf{s}}(s_i,s_j) \coloneqq d^{\pi}_{s_i,s_j}(s_i,s_j) = d^{\pi}(s_i,s_j). \tag{14}$$

The above result directly leads to that for any $\rho, \rho' \in \Delta(\mathcal{C}) \subseteq \Delta(\mathcal{S})$,

$$\mathcal{W}_{d^{\pi}}(\rho,\rho') = \mathcal{W}_{d^{\pi}_{\mathsf{s}}}(\rho,\rho'). \tag{15}$$

A.2 Additional Assumptions in the proof

Besides the key properties of the target CMDPs introduced in Assumption 4.1, we shall introduce some auxiliary assumptions for convenience as follows:

- **Bounded environment.** For the optimal policy $\pi^\star$ and any state $s \in \mathcal{S}$, without loss of generality, the minimum and maximum π**-bisimulation** distance between s and other states satisfies

$$\min_{s' \in \mathcal{S}} d_{\mathsf{s}}^\star(s, s') \mathbb{1}(d_{\mathsf{s}}^\star(s, s') > 0) \geq 1, \qquad \frac{D_{\mathsf{max}}}{2} \leq \max_{s' \in \mathcal{S}} d_{\mathsf{s}}^\star(s, s') \leq D_{\mathsf{max}}, \tag{16}$$

  for some universal positive constant D_{max}.

- **Deterministic environment.** Without loss of generality, we suppose the transition kernel is deterministic. Moreover, at any state $s \in \mathsf{supp}(\rho_{k+1}) \setminus \mathsf{supp}(\rho_k)$, there are a fixed portion of actions lead to the decreasing of the distance $\min_{s' \in \rho_k} d_{\mathsf{s}}^\star(s', s)$, i.e., for any transition sample (s_t, a_t, s_{t+1}),

$$\min_{s' \in \rho_k} d_{\mathsf{s}}^\star(s', s_{t+1}) = \begin{cases} \min_{s' \in \rho_k} d_{\mathsf{s}}^\star(s', s_t) - 1 & \text{with probability p} \\ \min_{s' \in \rho_k} d_{\mathsf{s}}^\star(s', s_t) + 1 & \text{otherwise.} \end{cases} \tag{17}$$

- **Closeness of the subsequent stages.** Starting from any state $s_0 = t \in \mathsf{supp}(\rho_{k+1}) \setminus \mathsf{supp}(\rho_k)$, if the sample trajectory shall enter $\mathcal{C}_{\mathsf{s}}(\rho_k, \pi_k^\star)$ before arriving at somewhere s obeying $d_{\mathsf{s}}^\star(s, t) = \max_{s \in \mathcal{S}} d_{\mathsf{s}}^\star(s, t)$. Then we suppose that the hitting point s is not far from the initial distribution ρ_k of the previous stage k, namely

$$\mathbb{E}_s[d_{\mathsf{s}}^\star(s, t)] \leq C_1 \min_{s' \in \rho_k} d_{\mathsf{s}}^\star(s', t), \tag{18}$$

  for some universal constant $C_1 > 0$.

We would like to note that these additional assumptions are useful for a concise proof to show the main intuition and idea of our proposed algorithm **GRADIENT**. These are some general assumptions without specific limitations. It is interesting to extend our main Theorem 4.1 to more general cases which shall be the further work.

A.3 Preliminaries: optimal transport and random walk

In this subsection, we introduce several key properties/facts of optimal transport, value function, and random walk, which play a crucial role in the proof of Theorem 4.1. The proofs for this subsection are deferred to Appendix A.5.

Preliminaries of optimal transport. We first introduce an important fact of the optimization problem in optimal transport. Suppose there exists two discrete sets $\mathcal{X}, \mathcal{Y}$ (possibly different). As it is well-known, for any two distributions $\mu \in \Delta(\mathcal{X}), \nu \in \Delta(\mathcal{Y})$, 1-Wasserstein distance with any metric d between μ and ν can be expressed as the optimal solution of the following linear programming (Kantorovich duality) problem:

$$\begin{aligned} &\max_{u,v} \sum_{x \in \mathcal{X}} u(x)\mu(x) - \sum_{y \in \mathcal{Y}} v(y)\nu(y), \\ &\text{subject to} \quad \forall x \in \mathcal{X}, y \in \mathcal{Y}: \ u(x) - v(y) \leq d(x, y), \\ &\qquad\qquad 0 \leq u \leq 1, 0 \leq v \leq 1. \end{aligned} \tag{19}$$

Note that, when $\mathcal{X} = \mathcal{Y} = \mathcal{S}$, the resulting claim in (19) is equivalent to [48, 61, 62].

Controlling performance gap. Inspired by the above problem formulation, the difference of the value function conditioned on two different states/state distribution can be controlled in the following lemma:

Lemma A.1 *For any $s, t \in \mathcal{S}$ and any policy π, we have*

$$|V_{\mathsf{s}}^\pi(s) - V_{\mathsf{s}}^\pi(t)| \leq d_{\mathsf{s}}^\pi(s, t). \tag{20}$$

In addition, for any two context distribution $\rho_1, \rho_2 \in \Delta(\mathcal{C}) \in \Delta(\mathcal{S})$, the performance gap associated with any policy π obeys

$$|V_{\mathsf{s}}^\pi(\rho_1) - V_{\mathsf{s}}^\pi(\rho_2)| \leq \mathcal{W}_{d^\pi}(\rho_1, \rho_2). \tag{21}$$

This fact connects the performance gap in RL and the Wasserstein distance between different state distribution/contextual distribution, taking advantage of the Kantorovich duality form of the optimal transport problem.

Definition and properties of random walk. Finally, we describe the following useful lemma which is essential in proving the main part of Theorem 4.1 when the starting state is not in the covering area of the optimal policy of the previous training stage

Lemma A.2 *The process $\{S_n : n \geq 1\}$ is called a random walk if $\{X_i : i \geq 1\}$ are iid Bernoulli distribution with $p < 1$ and $S_n = \sum_{i=1}^{n} X_i$. Starting from 0, the expectation of the* **stopping time** *N of hitting any $d > 0$ or $-D_{\mathsf{max}}$ and the probability p_α of hitting $d > 0$ obeys:*

$$\mathbb{E}[N] \leq \max\left\{ D_{\mathsf{max}} d \,,\, \sqrt{\frac{D_{\mathsf{max}}}{(1-2p) \wedge 1}} d \right\}, \quad p_\alpha \geq 1 - \frac{2d\sqrt{D_{\mathsf{max}}}}{d + D_{\mathsf{max}}}. \tag{22}$$

A.4 Proof of Theorem 4.1

With above preliminaries in hand, we are ready to embark on the proof for Theorem 4.1, which is divided into multiple steps as follows.

Step 1: decomposing the performance gap of interest. To begin with, we decompose the term of interest as follows:

$$\begin{aligned} V^{\pi^\star_{k+1}}(\rho_{k+1}) - V^{\pi^\star_k}(\rho_{k+1}) &\overset{\text{(i)}}{=} V_{\mathsf{s}}^{\pi^\star_{k+1}}(\rho_{k+1}) - V_{\mathsf{s}}^{\pi^\star_k}(\rho_{k+1}) \\ &\overset{\text{(ii)}}{=} V_{\mathsf{s}}^{\pi^\star}(\rho_{k+1}) - V_{\mathsf{s}}^{\pi^\star_k}(\rho_{k+1}) \\ &= V_{\mathsf{s}}^{\pi^\star}(\rho_{k+1}) - V_{\mathsf{s}}^{\pi^\star}(\rho_k) + V_{\mathsf{s}}^{\pi^\star}(\rho_k) - V_{\mathsf{s}}^{\pi^\star_k}(\rho_{k+1}) \\ &\overset{\text{(iii)}}{\leq} \mathcal{W}_{d^\star}(\rho_{k+1}, \rho_k) + V_{\mathsf{s}}^{\pi^\star_k}(\rho_k) - V_{\mathsf{s}}^{\pi^\star_k}(\rho_{k+1}), \end{aligned} \tag{23}$$

where (i) holds by the equivalence verified in (12b), (ii) follows from the value function of the optimal policy $\pi^\star$ (covering the entire state-action space) is the same as that of $\pi^\star_{k+1}$ in the region $\mathcal{C}(\rho_{k+1}, \pi^\star_{k+1})$, and (iii) comes from applying Lemma A.1 and the fact that $\pi^\star$ achieves the same value function as $\pi^\star_k$ in the region $\mathcal{C}(\rho_k, \pi^\star_k)$.

Step 2: controlling the second term in (23) in two cases. To proceed, it is observed that

$$\forall (s,t) \in \mathcal{C}_s(\rho_k, \pi^\star_k) \times \rho_{k+1} : \qquad V_{\mathsf{s}}^{\pi^\star_k}(s) - V_{\mathsf{s}}^{\pi^\star_k}(t) \leq d^\star_{\mathsf{s}}(s,t) \tag{24}$$

can directly leads to that $(u, v) = \left(V_{\mathsf{s}}^{\pi^\star_k}, V_{\mathsf{s}}^{\pi^\star_k}\right)$ is a feasible solution to the Wasserstein distance $\mathcal{W}_{d^\star}(\rho_k, \rho_{k+1})$ dual formulation in (19), which yields

$$\begin{aligned} V_{\mathsf{s}}^{\pi^\star_k}(\rho_k) - V_{\mathsf{s}}^{\pi^\star_k}(\rho_{k+1}) &= \sum_{s \in \mathcal{C}_s(\rho_k, \pi^\star_k)} V_{\mathsf{s}}^{\pi^\star_k}(s)\rho_k(s) - \sum_{s \in \mathcal{C}_s(\rho_{k+1}, \pi^\star_k)} V_{\mathsf{s}}^{\pi^\star_k}(s)\rho_{k+1}(s) \\ &\leq \mathcal{W}_{d^\star_{\mathsf{s}}}(\rho_k, \rho_{k+1}). \end{aligned} \tag{25}$$

As a result, we turn to show (24) instead of controlling $V_{\mathsf{s}}^{\pi^\star_k}(\rho_k) - V_{\mathsf{s}}^{\pi^\star_k}(\rho_{k+1})$. Towards this, we start from considering $V_{\mathsf{s}}^{\pi^\star_k}(s) - V_{\mathsf{s}}^{\pi^\star_k}(t)$ for any $s \in \rho_k$ and $t \in \rho_{k+1}$ in two different cases: (i) $t \in \mathcal{C}_{\mathsf{s}}(\rho_k, \pi^\star_k)$; (ii) otherwise.

In the first case when $t \in \mathcal{C}_{\mathsf{s}}(\rho_k, \pi^\star_k)$, since $\pi^\star_k$ is the same as the optimal policy $\pi^\star$ in $\mathcal{C}_{\mathsf{s}}(\rho_k, \pi^\star_k)$, invoking Lemma A.1 directly yields

$$V_{\mathsf{s}}^{\pi^\star_k}(s) - V_{\mathsf{s}}^{\pi^\star_k}(t) = V_{\mathsf{s}}^{\pi^\star}(s) - V_{\mathsf{s}}^{\pi^\star}(t) \leq d^\star_{\mathsf{s}}(s,t). \tag{26}$$

Step 3: focusing on case (ii). Then, we shall focus on the other case when

$$t \in \mathsf{supp}(\rho_{k+1}) \setminus \mathcal{C}_{\mathsf{s}}(\rho_k, \pi^\star_k). \tag{27}$$

Invoking Lemma A.2, without loss of generality, since in the stage k, we can't visit the region outside of $\mathcal{C}_{\mathsf{s}}(\rho_k, \pi_k^\star)$, we can define $\pi_k^\star(\cdot \mid s)$ to be uniformly random in the unseen region, i.e.,

$$\forall a \in \mathcal{A}, s \in \mathsf{supp}(\rho_{k+1}) \setminus \mathcal{C}_{\mathsf{s}}(\rho_k, \pi_k^\star): \qquad \pi_k^\star(a \mid s) = \frac{1}{A}. \tag{28}$$

Then we define p_{g} as the probability of the benign events B when we can hit some point s_{b} at the boundary of the region $\mathcal{C}_{\mathsf{s}}(\rho_k, \pi_k^\star)$ visited by the previous stage k before arriving at some limit point $s_{\mathsf{limit}} \coloneqq \arg\max_{s' \in \mathcal{S}} d_{\mathsf{s}}^\star(t, s')$. Invoking the assumption of the target CMDPs in (17), combind with the policy defined in (28), we observe that applying Lemma A.2 leads to

$$p_{\mathsf{g}} \geq 1 - \frac{2d_{\mathsf{s}}^\star(t, s_{\mathsf{b}})\sqrt{D_{\mathsf{max}}}}{d_{\mathsf{s}}^\star(t, s_{\mathsf{b}}) + D_{\mathsf{max}}} \geq 1 - \frac{2d_{\mathsf{s}}^\star(t, s_{\mathsf{b}})}{\sqrt{D_{\mathsf{max}}}} \geq 1 - \frac{2C_1 d_{\mathsf{s}}^\star(t, s)}{\sqrt{D_{\mathsf{max}}}}, \tag{29}$$

where the last inequality holds by the assumption in (16).

To continue, we express the value function at state t as

$$V_{\mathsf{s}}^{\pi_k^\star}(t) \geq p_{\mathsf{g}} \left(\mathbb{E}_B[\sum_{t=0}^{N} 0 \cdot \gamma^t] + \gamma^N V_{\mathsf{s}}^{\pi_k^\star}(s_{\mathsf{b}}) \right) - (1 - p_{\mathsf{g}}) = \gamma^N p_{\mathsf{g}} V_{\mathsf{s}}^{\pi_k^\star}(s_{\mathsf{b}}) - (1 - p_{\mathsf{g}}). \tag{30}$$

Similarly, we have

$$V_{\mathsf{s}}^{\pi_k^\star}(s) \leq (1 - p_{\mathsf{g}}) + p_{\mathsf{g}} \mathbb{E}[\sum_{t=0}^{N} (1-\gamma)\gamma^t + \gamma^N V_{\mathsf{s}}^{\pi_k^\star}(s_N)]. \tag{31}$$

We shall control the performance gap in two cases separately:

- When $V_{\mathsf{s}}^{\pi_k^\star}(s_{\mathsf{b}}) \geq V_{\mathsf{s}}^{\pi_k^\star}(s_N)$. we have

$$\begin{aligned} V_{\mathsf{s}}^{\pi_k^\star}(s) - V_{\mathsf{s}}^{\pi_k^\star}(t) &\leq V_{\mathsf{s}}^{\pi_k^\star}(s) - \gamma^N p_{\mathsf{g}} V_{\mathsf{s}}^{\pi_k^\star}(s') + (1 - p_{\mathsf{g}}) \\ &\leq p_{\mathsf{g}} \mathbb{E}[\sum_{t=0}^{N} (1-\gamma)\gamma^t + \gamma^N V_{\mathsf{s}}^{\pi_k^\star}(s_N)] - \gamma^N p_{\mathsf{g}} V_{\mathsf{s}}^{\pi_k^\star}(s_{\mathsf{b}}) + 2(1 - p_{\mathsf{g}}) \\ &\leq p_{\mathsf{g}} \mathbb{E}[\sum_{t=0}^{N} (1-\gamma)\gamma^t] + 2(1 - p_{\mathsf{g}}). \end{aligned} \tag{32}$$

- When $V_{\mathsf{s}}^{\pi_k^\star}(s_{\mathsf{b}}) < V_{\mathsf{s}}^{\pi_k^\star}(s_N)$, we have

$$\begin{aligned} V_{\mathsf{s}}^{\pi_k^\star}(s) - V_{\mathsf{s}}^{\pi_k^\star}(t) &\leq p_{\mathsf{g}} \mathbb{E}[\sum_{t=0}^{N} (1-\gamma)\gamma^t + \gamma^N V_{\mathsf{s}}^{\pi_k^\star}(s_N)] - \gamma^N p_{\mathsf{g}} V_{\mathsf{s}}^{\pi_k^\star}(s_{\mathsf{b}}) + 2(1 - p_{\mathsf{g}}) \\ &\leq p_{\mathsf{g}} \mathbb{E}[\sum_{t=0}^{N} (1-\gamma)\gamma^t + \gamma^N (N + V_{\mathsf{s}}^{\pi_k^\star}(s))] - \gamma^N p_{\mathsf{g}} V_{\mathsf{s}}^{\pi_k^\star}(s_{\mathsf{b}}) + 2(1 - p_{\mathsf{g}}) \\ &\leq p_{\mathsf{g}} \mathbb{E}\left[\sum_{t=0}^{N} (1-\gamma)\gamma^t + \gamma^N \big(N + V_{\mathsf{s}}^{\pi_k^\star}(s_{\mathsf{b}}) + (C_1 + 1) d^\star(s, t) \big) \right] \\ &\quad - \gamma^N p_{\mathsf{g}} V_{\mathsf{s}}^{\pi_k^\star}(s_{\mathsf{b}}) + 2(1 - p_{\mathsf{g}}) \\ &\leq p_{\mathsf{g}} \left[\sum_{t=0}^{N} (1-\gamma)\gamma^t + \gamma^N \big(N + (C_1 + 1) d^\star(s, t) \big) \right] + 2(1 - p_{\mathsf{g}}). \end{aligned} \tag{33}$$

Summing up the above two cases yields, for all $s \in \mathcal{C}(\rho_k, \pi_k^\star), t \in \mathsf{supp}(\rho_{k+1}) \setminus \mathcal{C}(\rho_k, \pi_k^\star)$,

$$
\begin{aligned}
&V_{\mathsf{s}}^{\pi_k^\star}(s) - V_{\mathsf{s}}^{\pi_k^\star}(t) \\
&\le p_{\mathsf{g}} \left[\sum_{t=0}^{N} (1-\gamma)\gamma^t + \gamma^N \big(N + (C_1+1) d^\star(s,t) \big) \right] + 2(1-p_{\mathsf{g}}) \\
&\le p_{\mathsf{g}} \mathbb{E}_B[N] + p_{\mathsf{g}}(C_1+1) d^\star(s,t) + 2(1-p_{\mathsf{g}}) \\
&\overset{\text{(i)}}{\le} \max\left\{ D_{\mathsf{max}} d^\star(s,t) \,,\, \sqrt{\frac{D_{\mathsf{max}}}{(1-2p)\wedge 1}} d^\star(s,t) \right\} + (C_1+1) d^\star(s,t) + \frac{4 d^\star(s,t)}{\sqrt{D_{\mathsf{max}}}} \\
&\le \left(\max\left\{ D_{\mathsf{max}} \,,\, \sqrt{\frac{D_{\mathsf{max}}}{(1-2p)\wedge 1}} \right\} + (C_1+1) + \frac{4}{\sqrt{D_{\mathsf{max}}}} \right) d^\star(s,t) \\
&\le c_{\mathsf{model}} d^\star(s,t)
\end{aligned} \tag{34}
$$

where (i) follows from Lemma A.2, and the last inequality holds by choosing

$$
c_{\mathsf{model}} \coloneqq \left(\max\left\{ D_{\mathsf{max}} \,,\, \sqrt{\frac{D_{\mathsf{max}}}{(1-2p)\wedge 1}} \right\} + (C_1+1) + \frac{4}{\sqrt{D_{\mathsf{max}}}} \right)
$$

as a large enough problem-dependent parameter.

Step 4: summing up the results. Combining the results (26) and (34) in two cases, we directly arrive at, for all $s \in \mathcal{C}(\rho_k, \pi_k^\star), t \in \mathsf{supp}(\rho_{k+1})$, the following fact is satisfied

$$
\begin{aligned}
&V_{\mathsf{s}}^{\pi_k^\star}(s) - V_{\mathsf{s}}^{\pi_k^\star}(t) \le c_{\mathsf{model}} d^\star(s,t), \\
&\to \frac{1}{c_{\mathsf{model}}} V_{\mathsf{s}}^{\pi_k^\star}(s) - \frac{1}{c_{\mathsf{model}}} V_{\mathsf{s}}^{\pi_k^\star}(t) \le d^\star(s,t).
\end{aligned} \tag{35}
$$

We observe that $(u, v) = \big(\frac{1}{c_{\mathsf{model}}} V_{\mathsf{s}}^{\pi_k^\star}, \frac{1}{c_{\mathsf{model}}} V_{\mathsf{s}}^{\pi_k^\star} \big)$ is a feasible solution to the Wasserstein distance $\mathcal{W}_{d_{\mathsf{s}}^\star}(\rho_k, \rho_{k+1})$ dual formulation in (19), we achieve

$$
V_{\mathsf{s}}^{\pi_k^\star}(\rho_k) - V_{\mathsf{s}}^{\pi_k^\star}(\rho_{k+1}) \le c_{\mathsf{model}} \mathcal{W}_{d_{\mathsf{s}}^\star}(\rho_k, \rho_{k+1}) \,. \tag{36}
$$

Finally, plugging (36) into (23) complete the proof by

$$
V^{\pi_{k+1}^\star}(\rho_{k+1}) - V^{\pi_k^\star}(\rho_{k+1}) \le (1 + c_{\mathsf{model}}) \mathcal{W}_{d^\star}(\rho_k, \rho_{k+1}) \,. \tag{37}
$$

A.5 Proof of auxiliary lemmas

Proof of Lemma A.1. As it is well known that $V_{\mathsf{s}}^\pi(s)$ is the fixed point of the Bellman operator $\mathcal{T}(V, \pi) \coloneqq \mathbb{E}_{a \sim \pi(\cdot \mid s)}[R(s,a,s) + \gamma \sum_{s'} P(s' \mid s, a, s) V]$. In addition, d_{s}^π is the fixed point of the operator [48]

$$
\begin{aligned}
\mathcal{F}(d)(s,t) \coloneqq &\big| \mathbb{E}_{a \sim \pi(\cdot \mid s)}[R(s,a,s) - \mathbb{E}_{a \sim \pi(\cdot \mid t)}[R(t,a,t) \big| \\
&+ \gamma \mathcal{W}_d \left(\mathbb{E}_{a \sim \pi(\cdot \mid t)}[P(\cdot \mid s,a,s)], \mathbb{E}_{a \sim \pi(\cdot \mid t)}[P(\cdot \mid t,a)] \right) .
\end{aligned} \tag{38}
$$

Armed with above facts, initializing $V_{\mathsf{s},0}^\pi = 0$ and $d_{\mathsf{s},0}^\pi = 0$, the update rules of $V_{\mathsf{s},n+1}^\pi$ and $d_{\mathsf{s},n+1}^\pi$ at the $(n+1)$-th iteration are defined as

$$
\begin{aligned}
V_{\mathsf{s},n+1}^\pi &= \mathcal{T}(V_{\mathsf{s},n}^\pi, \pi), \\
d_{\mathsf{s},n+1}^\pi &= \mathcal{F}(d_{\mathsf{s},n}^\pi)(s,t).
\end{aligned} \tag{39}
$$

With this in mind, for any $s, t \in \mathcal{S}$, we arrive at,

$$
\begin{aligned}
&\left|V^{\pi}_{\mathsf{s},n+1}(s) - V^{\pi}_{\mathsf{s},n+1}(t)\right| \\
&= \left| \mathbb{E}_{a\sim\pi(\cdot\,|\,s)}\left[R(s,a,s) + \gamma \sum_{s'} P(s'\,|\,s,a,s) V^{\pi}_{\mathsf{s},n}\right] \right. \\
&\qquad \left. - \mathbb{E}_{a\sim\pi(\cdot\,|\,t)}\left[R(t,a,t) + \gamma \sum_{s'} P(s'\,|\,t,a,t) V^{\pi}_{\mathsf{s},n}\right] \right| \\
&\le \left|\mathbb{E}_{a\sim\pi(\cdot\,|\,s)} R(s,a,s) - \mathbb{E}_{a\sim\pi(\cdot\,|\,t)} R(t,a,t)\right| + \gamma \left| \sum_{s'} P(s'\,|\,s,a,s) V^{\pi}_{\mathsf{s},n} - \sum_{s'} P(s'\,|\,t,a,t) V^{\pi}_{\mathsf{s},n} \right| \\
&\overset{\text{(i)}}{\le} \left|\mathbb{E}_{a\sim\pi(\cdot\,|\,s)} R(s,a,s) - \mathbb{E}_{a\sim\pi(\cdot\,|\,t)} R(t,a,t)\right| \\
&\qquad + \gamma \mathcal{W}_{d^{\pi}_{\mathsf{s},n}}\left(\mathbb{E}_{a\sim\pi(\cdot\,|\,s)}[P(s'\,|\,s,a,s)], \mathbb{E}_{a\sim\pi(\cdot\,|\,t)}[P(s'\,|\,t,a,t)]\right) \\
&= \mathcal{F}(d^{\pi}_{\mathsf{s},n})(s,t) = d^{\pi}_{\mathsf{s},n+1},
\end{aligned} \tag{40}
$$

where (i) holds by $(u, v) = \left(V^{\pi}_{\mathsf{s},n}, V^{\pi}_{\mathsf{s},n}\right)$ is a feasible solution to the Wasserstein distance $\mathcal{W}_{d^{\pi}_{\mathsf{s},n}}\left(\mathbb{E}_{a\sim\pi(\cdot\,|\,s)}[P(s'\,|\,s,a,s)], \mathbb{E}_{a\sim\pi(\cdot\,|\,t)}[P(s'\,|\,t,a,t)]\right)$ dual formulation in (19).

As a result, we have

$$
\forall n = 1, 2, \cdots: \qquad \left|V^{\pi}_{\mathsf{s},n+1}(s) - V^{\pi}_{\mathsf{s},n+1}(t)\right| \le d^{\pi}_{\mathsf{s},n+1}, \tag{41}
$$

which directly yields

$$
\forall s, t \in \mathcal{S}: \qquad |V^{\pi}_{\mathsf{s}}(s) - V^{\pi}_{\mathsf{s}}(t)| \le d^{\pi}_{\mathsf{s}}(s,t). \tag{42}
$$

The above fact indicates that for any two distribution $\rho, \rho' \in \Delta(\mathcal{S})$,

$$
V^{\pi}_{\mathsf{s}}(\rho_k) - V^{\pi}_{\mathsf{s}}(\rho_{k+1}) = \sum_{s\in\mathcal{S}} V^{\pi}_{\mathsf{s}}(s)\rho(s) - \sum_{s\in\mathcal{S}} V^{\pi}_{\mathsf{s}}(s)\rho'(s) \le \mathcal{W}_{d^{\pi}_{\mathsf{s}}}(\rho,\rho') = \mathcal{W}_{d^{\pi}}(\rho,\rho'), \tag{43}
$$

where the penultimate inequality arises from $(u, v) = \left(V^{\pi}_{\mathsf{s}}, V^{\pi}_{\mathsf{s}}\right)$ is a feasible solution to the Wasserstein distance $\mathcal{W}_{d^{\pi}_{\mathsf{s}}}(\rho,\rho')$ dual formulation in (19), and the last equality holds by (15).

Proof of Lemma A.2 Before starting, we denote p_α as the probability that stopping at d and N as the stopping time. Then we consider the terms of interest in two cases. In the first case when $p = \frac{1}{2}$, invoking the facts in [63] directly leads to

$$
N = D_{\mathsf{max}} d, \qquad p_\alpha = 1 - \frac{d}{d + D_{\mathsf{max}}}. \tag{44}
$$

So the remainder of the proof will focus on the other case when $p \neq \frac{1}{2}$. Following the proof in [63], by basic calculus, it is easily verified that the following two centered terms are all martingales

$$
S_n - (2p-1)n, \qquad S_n^2 + 2(1-2p)S_n + (2p-1)^2 n^2 + 4p(p-1)n. \tag{45}
$$

In continue, it can also be verified that

$$
\mathbb{E}[S_N - (2p-1)N] = p_\alpha d - (1-p_\alpha) D_{\mathsf{max}} = 0, \tag{46}
$$

which leads to

$$
p_\alpha = \frac{D_{\mathsf{max}} + (2p-1)N}{d + D_{\mathsf{max}}}. \tag{47}
$$

Observing that

$$
\mathbb{E}[S_n^2 + 2(1-2p)S_n + (2p-1)^2 n^2 + 4p(p-1)n] = 0 \tag{48}
$$

yields

$$
(2p-1)^2 N^2 + [(D_{\mathsf{max}} - d)(1-2p) - (sp-1)^2 - 1]N + D_{\mathsf{max}} d = 0, \tag{49}
$$

and then implies

$$N \leq \sqrt{\frac{D_{\mathsf{max}} d}{1-2p}} \leq \sqrt{\frac{D_{\mathsf{max}}}{1-2p}} d. \tag{50}$$

Plugging in (50) into (47) leads to

$$p_\alpha \geq \frac{D_{\mathsf{max}} - \sqrt{D_{\mathsf{max}} d}}{d + D_{\mathsf{max}}} = 1 - \frac{\sqrt{D_{\mathsf{max}} d} + d}{d + D_{\mathsf{max}}} \geq 1 - \frac{2d\sqrt{D_{\mathsf{max}}}}{d + D_{\mathsf{max}}}, \tag{51}$$

which follows from (16) in Appendix A.2.

Finally, summing up the two cases, we arrive at

$$N \leq \max\left\{ D_{\mathsf{max}} d \,,\, \sqrt{\frac{D_{\mathsf{max}}}{(1-2p) \wedge 1}} d \right\}, p_\alpha \geq 1 - \frac{2d\sqrt{D_{\mathsf{max}}}}{d + D_{\mathsf{max}}}. \tag{52}$$

A.6 Table of notations

We summarize useful notations in the proof here.

Table 1: Notations

Symbol	Definition
$\Delta(X)$	the set of probability measures on X
$\mathsf{supp}(\rho)$	support of distribution ρ
$o^\pi(s;\rho)$	$(1-\gamma)\sum_{t=0}^\infty \gamma^t \mathbb{P}\big(s_t = s \,|\, c \sim \rho; \pi\big)$
$o^\pi(s,a;\rho)$	$o^\pi(s;\rho)\pi(a \,|\, s)$
$\mathcal{C}_s(\rho,\pi)$	$\{s : o^\pi(s;\rho) > 0\}$
$\mathcal{C}(\rho,\pi)$	$\{(s,a) : o^\pi(s,a;\rho) > 0\}$
$V_{\mathsf{s}}^\pi(s)$	$\mathbb{E}\left[\sum_{t=0}^\infty \gamma^t R(s_t,a_t,c') \,\middle|\, s_0 = s, c = s; \pi\right]$
$d_{\mathsf{s}}^\star$	bisimulation distance $d_{\mathsf{s}}^{\pi^*}$ under the optimal policy π^*
$a \wedge b$	$\min(a,b)$

B Experiment Details

In this section, we discuss details that could not be included in the main paper due to space limitations. This includes hyperaparameters of the algorithms, additional details about the environment, and visualizations that assist the qualitative analysis. The experiments were conducted on a desktop computer with an Intel Core i7-8700K CPU @ 3.70GHz 12-Core Processor, an Nvidia RTX 2080Ti graphics card and 64GB of RAM. All evaluation results are based on 30 episodes over 3 random seeds.

B.1 Maze

We show the observation, action and context space in Table 2. The state, action and context space are discrete. For this toy-like example, we fix the maze layout throughout. The context is the initial position. The observation is the flattened value representation of the maze, including the goal, the current position of the agent, and the layout.

Table 2: Maze Environment Specifications

Dim.	Discrete Observation Space
0-120	Cell Type: $\{0:\text{free}, 1:\text{wall}, 2:\text{agent}, 3:\text{goal}\}$
Index	Discrete Action Space
0	Go north
1	Go south
2	Go west
3	Go east
Index	Discrete Context Space
0-50	Initial position

Table 3: Maze Hyperparameters

PPO Learner	value
gamma	0.99
learning_rate	0.0001
n_steps	100
ent_coef	0.1
total_timesteps	100000
GRADIENT hyperparameters	value
reward threshold $\bar{G}$	-15

For this task, we use the PPO implementation in the `StableBaseline3` library [54]. We list the hyperparameters in the Table 3 (set to the default value if not mentioned in the table). We conduct hyperparameter grid search mainly on the following hyperparameters: $(learning_rate, ent_coef, n_steps) \in \{0.0001, 0.0003, 0.001\} \times \{0.001, 0.01, 0.1\} \times \{10, 100, 200\}$

We visualize the intermediate task distributions generated by GRADIENT in Figure 8 with a $\Delta\alpha = 0.05$. For $\Delta\alpha = 0.1$ and 0.2 are in the same figure as well (just with different gaps between two consecutive α). The π-contextual-distance is computed using the A* path finding algorithm. Although it is not always realistic to have access to the optimal policy, this toy example is just to show the effectiveness of an appropriate contextual distance metric. More choices of the contextual distance could be explored in the future work.

B.2 PointMass

In PointMass, the agent needs to navigate a pointmass through a wall with a small gap at an off-center position to reach the goal. The context is a 2-dimension vector representing the width and position of the gap on the wall. Following the setting in the original SPDL paper [8], the target distribution is an isotropic Gaussian distribution centered at $[2.5, 0.5]$ with a negligible variance $[4\times10^{-3}, 3.75\times10^{-3}]$ (which is effectively a point as shown in Figure 5a). The source distribution is an isotropic Gaussian distribution centered at $[0, 4.25]$ with a variance of $[2, 1.875]$. We show the observation, action and context space in Table 4.

For the baseline implantation, we use and modify the code base provided in [8]. We also use the best hyperparameters found by [8]. For ALP-GMM, they conduct hyperparameter grid search over $(p_{\text{RAND}}, n_{\text{ROLLOUT}}, s_{\text{BUFFER}}) \in \{0.1, 0.2, 0.3\} \times \{50, 100, 200\} \times \{500, 1000, 2000\}$. For Goal-GAN, they conduct grid search over $(\delta_{\text{NOISE}}, n_{\text{ROLLOUT}}, p_{\text{SUCCESS}}) \in \{0.025, 0.05, 0.1\} \times \{50, 100, 200\} \times \{0.1, 0.2, 0.3\}$. Due to unknown issue, the scale of episodic reward we get (about 0-100) is different from what is shown in [8] (0-10), nevertheless the trend of the training curves is

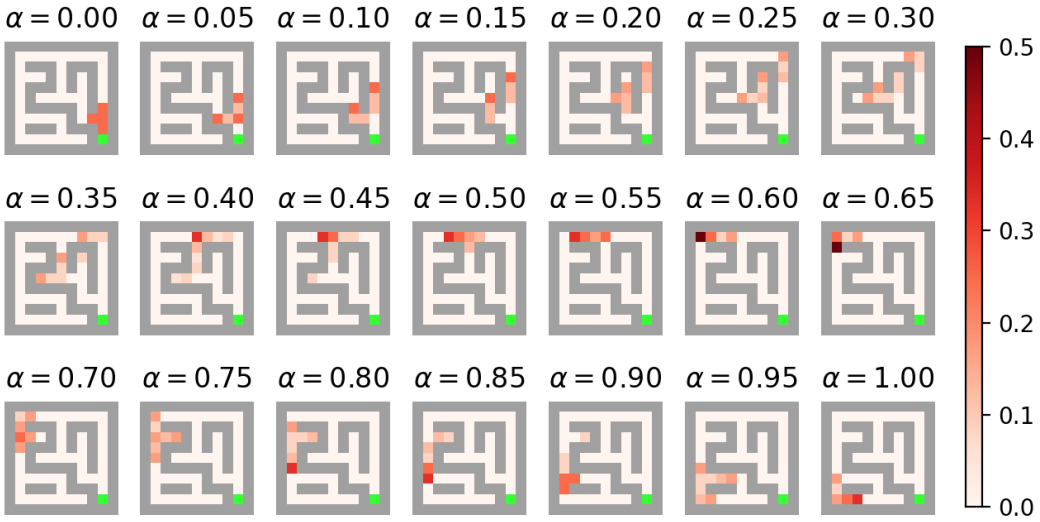

Figure 8: Intermediate task distributions generated by GRADIENT.

very similar. We visualize the context distribution of ALP-GMM and Goal-GAN in Figure 9. From this figure, we can reason why they underform GRADIENT and SPDL: since they cannot specify the target distribution, it is very difficult for them to learn to navigate through a narrow door at a specific position.

Table 4: PointMass Environment Specifications

Dim.	Continuous Observation Space	range
0	x position	$[-4, 4]$
1	x velocity	$[-\inf, \inf]$
2	y position	$[-4, 4]$
3	y velocity	$[-\inf, \inf]$
Dim.	Continuous Action Space	
0	x force	$[-10, 10]$
1	y force	$[-10, 10]$
Dim.	Continuous Context Space	
0	gate position	$[-4, 4]$
1	gate width	$[0.5, 8]$

B.3 FetchPush

In FetchPush [56], the objective is to use the gripper to push the box to a goal position. The observation space is a 28-dimension vector, including information about the goal. The context is a 2-dimension vector representing the goal position on a surface. The target distribution is a uniform distribution over the circumference of a half-circle (Figure 5b). The source distribution is a uniform distribution over a square region centered at the box position, excluding the region within a certain radius of the object. We use this experiment to highlight the importance of the capability to handle arbitrary distributions rather than only the parametric Gaussian distributions. Since `SPDL` can only deal with parametric Gaussian distribution, we first fit the target and the source distribution with two Gaussian distributions and feed them into the baseline algorithms. The source Gaussian is $[1.14655655, 0.74819359]$ with a variance array of $[[0.0141083, 0.00055327], [0.00055327, 0.0149638]]$, and the target Gaussian is $[1.33561676, 0.74819359]$ with a variance array of $[[8.89519060e-03, -1.34467507e-18], [-1.34467507e-18, 4.59090909e-02]]$. We visualize the context distribution of ALP-GMM and Goal-GAN in Figure 10.

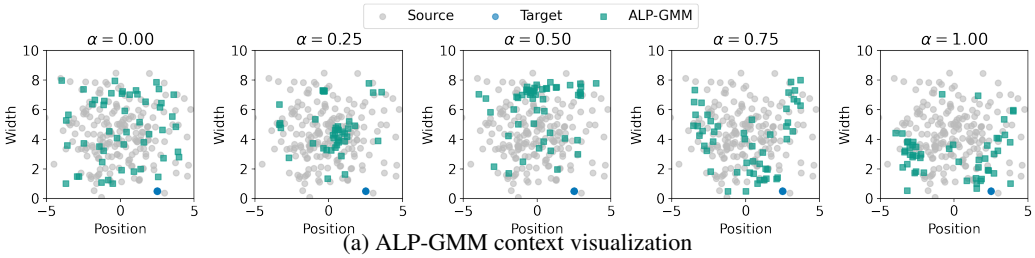

(a) ALP-GMM context visualization

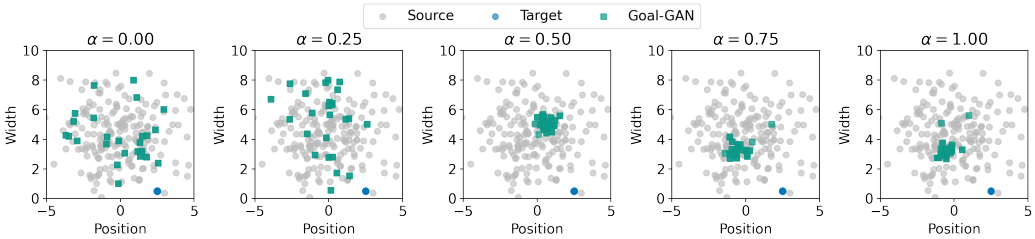

(b) Goal-GAN context visualization

Figure 9: Visualizations of context distributions and curricula in PointMass. The contexts are taken from environment steps at 10k, 50k, 100k, 200k, and 300k.

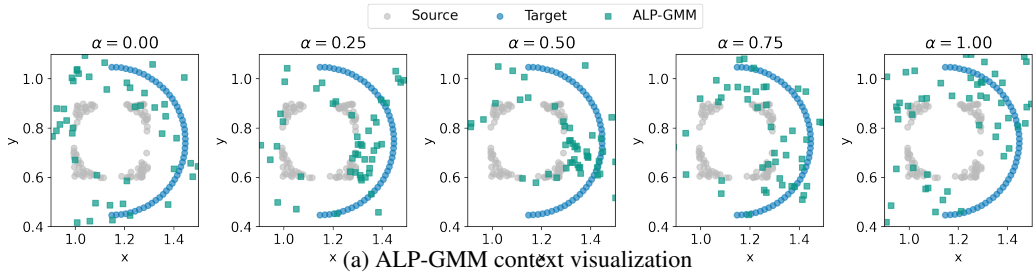

(a) ALP-GMM context visualization

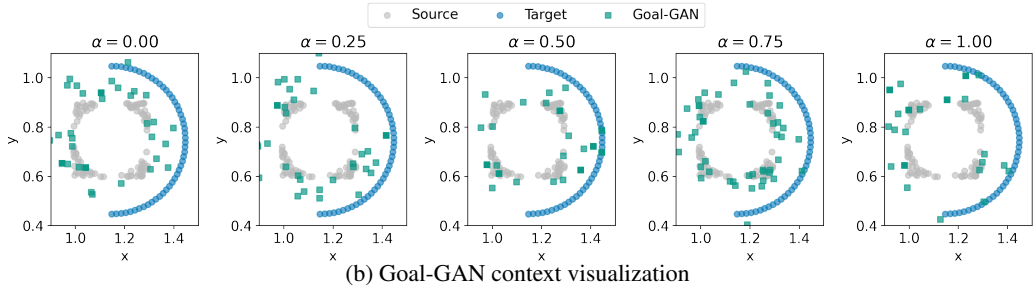

(b) Goal-GAN context visualization

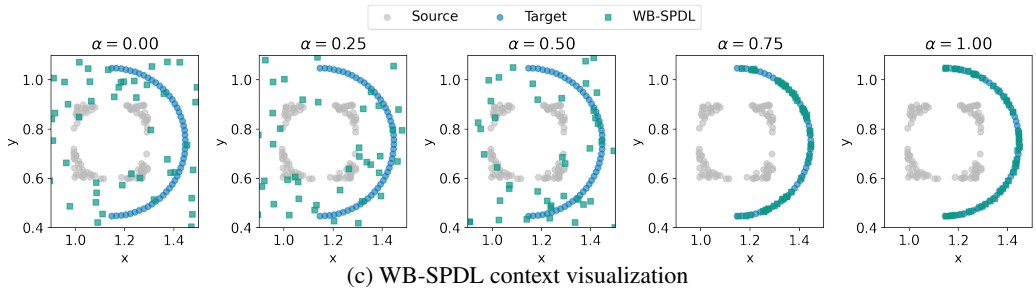

(c) WB-SPDL context visualization

Figure 10: Visualizations of context distributions and curricula in FetchPush. The contexts are taken from environment steps at 5k, 50k, 100k, 250k, and 500k for ALP-GMM and Goal-GAN, 100, 500, 1000, 1200, 1500 for WB-SPDL.

Table 5: PointMass Hyperparameters

SAC Learner	value
train_freq	5
buffer_size	10000
gamma	0.95
learning_rate	0.0003
learning_starts	500
batch_size	64
net architecture	[64, 64]
activation_fn	Tanh
SPDL hyperparameter	**value**
α offset	25 iterations
ζ	1.1
KL threshold	8000
maximum KL	0.05
steps per iteration	2048
ALP-GMM hyperparameter	**value**
random task ratio	0.2
fit rate (number of episodes between two fit of GMM)	200
max size (maximal number of episodes for computing ALP)	1000
GOAL-GAN hyperparameter	**value**
state noise level	0.05
fit rate (number of episodes between two fit of GAN)	25
probability to sample state with noise	0.05
GRADIENT hyperparameter	**value**
$\Delta\alpha$	0.1
Reward threshold $\bar{G}$	40

Table 6: FetchPush Environment Specifications

Dim.	Continuous Observation Space	range
0-2	grip_pos	$[-\inf, \inf]$
3-5	object_pos	$[-\inf, \inf]$
6-8	object_rel_pos	$[-\inf, \inf]$
9-10	gripper_state	$[-\inf, \inf]$
11-13	object_rot	$[-\inf, \inf]$
14-16	object_velp	$[-\inf, \inf]$
17-19	object_velr	$[-\inf, \inf]$
20-22	grip_velp	$[-\inf, \inf]$
23-24	gripper_vel	$[-\inf, \inf]$
25-27	goal_pos	$[-\inf, \inf]$
Index	Continuous Action Space	
0-2	pos_ctrl	$[-1, 1]$
3	gripper_ctrl	$[-1, 1]$
Dim.	Continuous Context Space	
0	goal_x_pos	$[0.8, 1.5]$
1	goal_y_pos	$[0.4, 1.1]$

Table 7: FetchPush Hyperparameters

SAC Learner	value
train_freq	1
buffer_size	100000
gamma	0.99
learning_rate	0.001
learning_starts	1000
batch_size	256
net architecture	[64, 64, 64]
activation_fn	Tanh
SPDL hyperparameter	**value**
α offset	0 iterations
ζ	1.0
KL threshold	20
maximum KL	0.05
steps per iteration	5000
ALP-GMM hyperparameter	**value**
random task ratio	0.3
fit rate (number of episodes between two fit of GMM)	200
max size (maximal number of episodes for computing ALP)	1000
GOAL-GAN hyperparameter	**value**
state noise level	0.1
fit rate (number of episodes between two fit of GAN)	200
probability to sample state with noise	0.3
GRADIENT hyperparameter	**value**
$\Delta\alpha$	0.2
Reward threshold $\bar{G}$	-25

C Algorithm Details

C.1 Exact Computation of π-contextual-distance for Maze

For the maze, we compute the exact π-contextual-distance using dynamic programming. We iteratively compute the π-contextual-distance (represented by a matrix) until the maximum difference in metric estimate between successive iterations is smaller than a tolerance. Due the Assumption 4.1, the computation of π-contextual-distance is largely simplified (since the context space is essentially overlapped with the state space). We present the pseudocode as follows:

C.2 Learning Embeddings to Encode Non-Euclidean Distance

Due to the computation complexity of the free-support Wasserstein barycenters, it is generally difficult to compute them, especially for non-euclidean cost metric. There are some attempts to use neural networks and stochastic gradient descent to solve for the barycenters approximately [58]. Another possible route is to learn embeddings such that the euclidean interpolation in the latent space assembles the interpolation in the non-euclidean original space. In the fixed-support setting, Deep Wasserstein Embedding (DWE) [59] uses siamese networks to learn an latent space where the euclidean distance approximates the Wasserstein distance in the original space. We could adopt a similar method but for the free-support, i.e., continuous context space.

Let the encoder be $\epsilon(\cdot)$ and decoder be $\delta(\cdot)$, given pairs of contexts $\{c_1^i, c_2^i\}_{i=1,\dots,N}$ and their contextual distance $\{d(c_1^i, c_2^i)\}_{i=1,\dots,N}$, the global objective funtion is to minimize

$$\min_{\epsilon,\delta} \sum_i \left\| \|\epsilon(c_1^i) - \epsilon(c_2^i)\|^2 - d(c_1^i, c_2^i) \right\|^2 + \lambda \sum_{k=1,2} \|\delta(\epsilon(c_k^i)) - c_k^i\|^2 \tag{53}$$

The first term is the loss for distance embedding, and the second term is for reconstruction. Other regularization could be added to improve the robustness. After obtaining the trained encoder and decoder, we can first encode the source and target contexts to the latent space, compute the Wasserstein barycenter in the latent space and finally decode the latent barycenter back to the original context space. The main algorithm GRADIENT with distance embedding is shown in Algorithm 3.

Algorithm 3: GRADIENT with Distance Embedding

Input: Source task distribution $\mu(c)$, target task distribution $\nu(c)$, interpolation factor $\Delta\alpha$, reward threshold $\bar{G}$.

Initialize the agent policy π;

$\alpha \leftarrow 0$;

for k *in* $0, 1, 2, ..., K$ **do**

  if $k == 0$ **then**

    $\rho(c) \leftarrow \mu(c)$;

  else

    $\rho(c) \leftarrow \delta(\texttt{ComputeBarycenter}(\epsilon(\mu), \epsilon(\nu), \alpha, l_2))$;

  $\{c_i, R_i\}_{i=1,\dots,M} \leftarrow$ Optimize π in the task distribution $\rho(c)$ until $G > \bar{G}$ (potentially with some exploration noise for c), and return recent M sampled context c_i and corresponding episodic rewards R_i;

  Estimate $J^\pi(c)$ from $\{c_i, R_i\}_{i=1,\dots,M}$ using Gaussian Process;

  `// Use the absolute difference between episodic reward to define the distance metric as an example`

  Define $d^\pi(c_i, c_j) := |J^\pi(c_i) - J^\pi(c_j)|$;

  Encoder ϵ, Decoder $\delta \leftarrow \texttt{EmbedDistanceMetric}(d^\pi)$; `// Algorithm 4`

  $\alpha \leftarrow \alpha + \Delta\alpha$;

  $\mu(c) \leftarrow \rho(c)$;

Output: Agent policy π

Algorithm 4: EmbedDistanceMetric

Input: contextual distance metric d.

Initialize encoder ϵ and decoder δ;

Uniformly sample pairs of contexts from the context range $\{c_1^i, c_2^i\}_{i=1,\dots,N}$; `// This step is cheap since there is no interaction with environment required.`

Compute $\{d(c_1^i, c_2^i)\}_{i=1,\dots,N}$;

Train encoder and decoder by minimizing (53);

Output: Encoder ϵ, Decoder δ

We show an example of using Algorithm 3. Here we consider a classical U-shaped maze with continuous state, action and context space. We assume that the source and target distributions are two Gaussian distributions at the two ends of the maze. Due to the existence of the obstacle in the middle, it is not appropriate to use the l_2 distance as the contextual distance. In this case, we use $d^\pi(c_1, c_2) := |J^\pi(c_1) - J^\pi(c_2)|$.

The interpolation results are shown in Figure 11. At the first a few stages, since the agent does not have a good policy as well as a good estimate of the episodic reward, the interpolation results are not very good. However, since α is small, the barycenters are relatively close to the source distribution, so GRADIENT does not generate too unreasonable tasks. With the learning progressing, the estimates become better and better and therefore produces much better interpolation results.

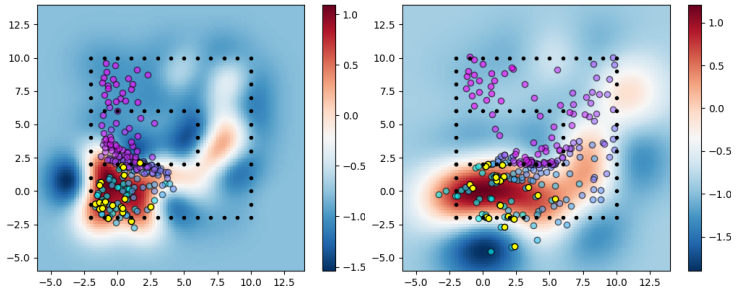

(a) Stage 1

(b) Stage 2

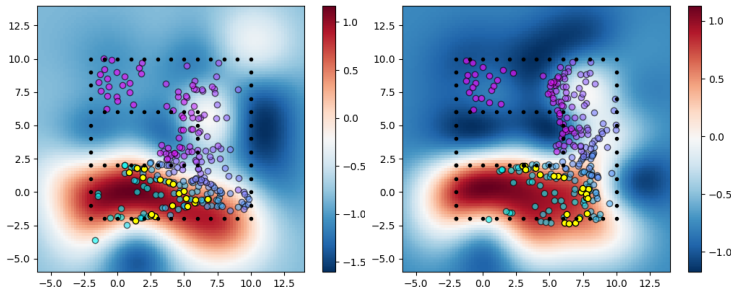

(c) Stage 3

(d) Stage 4

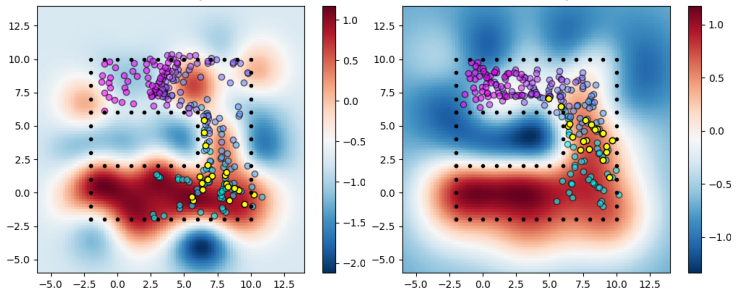

(e) Stage 5

(f) Stage 6

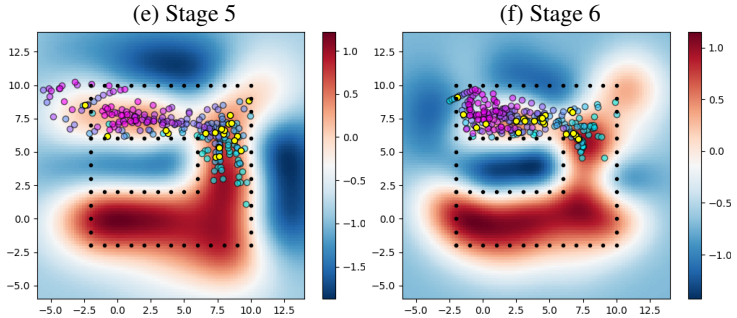

(g) Stage 7

(h) Stage 8

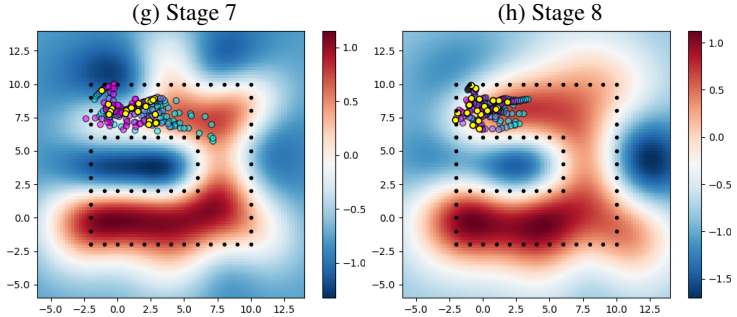

(i) Stage 9

(j) Stage 10

Figure 11: U-Maze interpolation results. The original source and target distributions are two Gaussian centered at $[0, 0]$ and $[0, 8]$. The color of the heat map represents the estimate of $J^\pi(c)$. The colored circle represent the Wasserstein interpolation. More specifically, the circles with cyan to purple color represent the interpolation results from the current source to target with $\alpha = [0, 1]$. The yellow circles highlight the barycenter corresponding to the current stage.