

# PAnDR: Fast Adaptation to New Environments from Offline Experiences via Decoupling Policy and Environment Representations

Tong Sang<sup>1\*</sup>, Hongyao Tang<sup>1\*</sup>, Yi Ma<sup>1</sup>, Jianye Hao<sup>1†</sup>, Yan Zheng<sup>1†</sup>, Zhaopeng Meng<sup>1</sup>, Boyan Li<sup>1</sup> and Zhen Wang<sup>2</sup>

<sup>1</sup>College of Intelligence and Computing, Tianjin University

<sup>2</sup>School of Artificial Intelligence, Optics and Electronics (iOPEN) and School of Cybersecurity, Northwestern Polytechnical University

{2019218044, bluecontra, mayi, jianye.hao, yanzheng, mengzp, liboyan}@tju.edu.cn, w-zhen@nwpu.edu.cn

## Abstract

Deep Reinforcement Learning (DRL) has been a promising solution to many complex decision-making problems. Nevertheless, the notorious weakness in generalization among environments prevent widespread application of DRL agents in real-world scenarios. Although advances have been made recently, most prior works assume sufficient online interaction on training environments, which can be costly in practical cases. To this end, we focus on an *offline-training-online-adaptation* setting, in which the agent first learns from offline experiences collected in environments with different dynamics and then performs online policy adaptation in environments with new dynamics. In this paper, we propose Policy Adaptation with Decoupled Representations (PAnDR) for fast policy adaptation. In offline training phase, the environment representation and policy representation are learned through contrastive learning and policy recovery, respectively. The representations are further refined by mutual information optimization to make them more decoupled and complete. With learned representations, a *Policy-Dynamics Value Function* (PDVF) [Raileanu *et al.*, 2020] network is trained to approximate the values for different combinations of policies and environments from offline experiences. In online adaptation phase, with the environment context inferred from few experiences collected in new environments, the policy is optimized by gradient ascent with respect to the PDVF. Our experiments show that PAnDR outperforms existing algorithms in several representative policy adaptation problems.

## 1 Introduction

Reinforcement learning (RL) has achieved great successes in many fields, e.g., Game Playing [Badia *et al.*, 2020], Software Testing [Zheng *et al.*, 2019; Zheng *et al.*, 2021], Robotics [Andrychowicz *et al.*, 2020], Industrial Supply

Chain [Ma *et al.*, 2021] and so on. Despite of the successes, conventional RL is well known to be lack of generalization ability [Packer *et al.*, 2018]. To be specific, a policy well optimized in a specific environment can be catastrophically poor even in similar environments. The poor generalization among environments and policies severely prevents more practical applications of RL algorithms.

In the literature on RL generalization [Kirk *et al.*, 2021], many approaches like data augmentation [Yarats *et al.*, 2021] and domain randomization [Peng *et al.*, 2018] are proposed to increase the similarity between training environments and the testing environments (i.e., the ones to generalize or adapt to). These methods often require knowing the variations of environment and having the access to environment generation. In contrast, some other works aim to achieve fast adaptation in testing environments without the need for such access and knowledge. Gradient-based Meta RL, e.g., MAML [Finn *et al.*, 2017], aims to learn a meta policy which can adapt to a testing environment within few policy gradient steps. Also in this branch, context-based Meta RL, e.g., PEARL [Rakelly *et al.*, 2019], leverages a context-conditioned policy which allows to achieve adaptation by the generalization among contexts. Useful context representation is learned in training environments to capture the variations; then the context of a testing environment can be inferred from few probed interaction experiences [Fu *et al.*, 2021].

Most prior works consider an *online* training setting where the agent is allowed to interact with training environments arbitrarily within the interaction budget. However, in real-world problems, online interaction is usually expensive while offline experiences are often available and relatively sufficient. For example, in E-commerce, adapting the advertising strategies obtained from massive offline ads-users interaction data to a new user can effectively save the budgets of advertisers. On this point, a recent work [Raileanu *et al.*, 2020] studies fast adaptation under an *offline-training-online-adaptation* setting, where only offline experiences collected on training environments are available before online policy adaptation in new environments. This is a more practical and general setting which is meanwhile more challenging, since the policy learning and variation capture need to be conducted in an offline manner. To solve this problem, Raileanu *et al.* [2020] propose an extended value function, called Policy-Dynamic Value Network (PDVF). By definition, a PDVF ad-

\*Equal contribution.

†Jianye Hao and Yan Zheng are the corresponding authors.

ditionally takes an environment representation and an policy representation as inputs, and evaluates the corresponding values. In principle, the explicit representation allows the generalization of value estimation among environments and policies, which can be utilized to realize policy adaptation. To be specific, given the offline experiences generated from different combinations of training policies and training environments, the representations of environment and policy are learned from transitions and state-action sequences by dynamics prediction and policy recovery respectively. Thereafter, a quadratic form of PDVF is approximated, according to which the policy is optimized during online adaptation with the environment representation inferred from few interactions in the testing environment. The first drawback is that the representations of environment and policy can have redundant and entangled information since they are from the same interaction experiences of the corresponding combination. It negatively affects the generalization conveyed by the representations which is vital to fast adaptation. Another limitation is that the quadratic approximation of PDVF inevitably makes a sacrifice in expressivity for convenient closed-form optimization, in turn this cripples the optimality of policy obtained during adaptation.

In this paper, we focus on the offline-training-online-adaptation setting and follow the paradigm of [Raileanu *et al.*, 2020]. We propose **Policy Adaptation with Decoupled Representations (PAnDR)** for more effective fast adaptation of RL. The conceptual illustration of PAnDR is shown in Fig. 1. Compared to transition-level dynamics prediction adopted by [Raileanu *et al.*, 2020], we adopt trajectory-level context contrast for environment representation learning. By this means, the variation is captured in an integral manner and the inter-environment difference is emphasized, thus offering more useful information and better robustness. One can link our comparison here to that between reconstruction and contrastive learning, where the superiority of the latter has been widely demonstrated [He *et al.*, 2020; Laskin *et al.*, 2020; Fu *et al.*, 2021]. Moreover, we propose an information-theoretic representation learning method to refine environment and policy representations for information decoupleness and completeness. This is achieved by mutual information minimization and maximization among environment and policy representations and their joint. In addition, distinct to the quadratic form in [Raileanu *et al.*, 2020], PAnDR leverages a typical multi-layer nonlinear neural network for the approximation of the PDVF. Then during online adaptation, the policy is optimized by gradient ascent. In this sense, PAnDR gets rid of such constraints to improve value approximation and policy adaptation. In experiments, we evaluate PAnDR in a series of fast adaptation tasks and analyze the contribution of each part of proposed methods.

We summarize our main contributions as follows: 1) We leverage trajectory-level contrastive learning and propose a novel mutual information based approach for better environment and policy representations. 2) We get rid of the constraint of prior quadratic approximation of PDVF and demonstrate the effectiveness of nonlinear approximation in improving online policy adaptation. 3) We show that PAnDR consistently outperforms representative existing algorithms in our experiments, sometimes with a large margin.

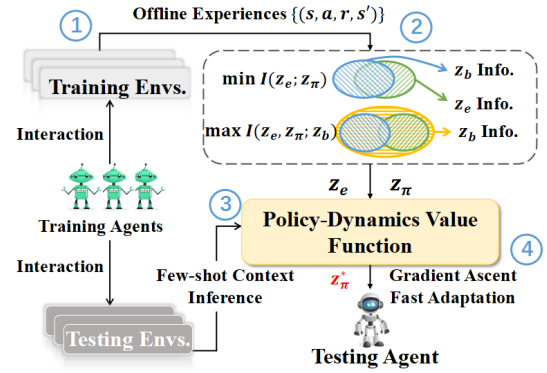


Figure 1: A conceptual illustration of PAnDR for fast adaptation with offline experiences (best view in numbers). With offline experiences collected by training policies and environments (①), the representations for environment ( $z_e$ ) and policy ( $z_\pi$ ) are learned, which are further improved by mutual information optimization with extra joint representation  $z_b$  (②). Conditioning on learned representations, a PDVF is trained. For adaptation, the representation of testing environment is inferred with few-shot online context (③), based on which the policy (representation) is adapted according to PDVF (④).

## 2 Preliminaries

First, we briefly introduce the necessary preliminaries. A detailed discussion on related works is provided in Appendix D.

### 2.1 Reinforcement Learning

Consider a Markov Decision Process (MDP) defined as  $\mathcal{M} = \langle S, A, P, R, \gamma \rangle$ , with the state space  $S$ , the action space  $A$ , the transition probability  $P : S \times A \times S \rightarrow [0, 1]$ , the reward function  $R : S \times A \rightarrow \mathbb{R}$  and the discount factor  $\gamma \in [0, 1]$ . An agent interacts with the MDP  $\mathcal{M}$  by performing its policy  $\pi : S \times A \rightarrow [0, 1]$ , generating the trajectories as  $a_t \sim \pi(\cdot | s_t)$ ,  $s_{t+1} \sim P(\cdot | s_t, a_t)$  and  $r_t = R(s_t, a_t)$ . The value function of policy  $\pi$  is defined as  $V^\pi(s) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s]$ . The objective of an RL agent is to optimize its policy to maximize the expected discounted cumulative rewards  $J(\pi) = \mathbb{E}_{s_0 \sim \rho_0(\cdot)}[V^\pi(s_0)]$ , where  $\rho_0$  is the initial state distribution.

### 2.2 Mutual Information Optimization

Mutual Information (MI) is a widely used metric that estimates the relationship between pairs of variables in machine learning. Given two random variables  $X$  and  $Y$ , MI is a measure of dependency among them, defined as  $I(X; Y) = \mathbb{E}_{p(x,y)} \left[ \frac{p(x,y)}{p(x)p(y)} \right]$ , with the joint probability distribution  $p(x, y)$  and the marginal probability distributions  $p(x), p(y)$ . MI estimation and optimization are fundamental techniques in representation learning, with Information Bottleneck [Alemi *et al.*, 2017] as a representative.

In most problems, it is intractable to compute MI directly since the exact distributions of interested variables are not known. Thus, approximate estimation of MI is usually resorted to, which is non-trivial especially in high-dimensional space. For the general purpose of MI optimization, lower-bound and upper-bound approximations of MI are often used as surrogate objectives [Poole *et al.*, 2019]. Among the options, the variational approximation is a common way to com-

pute lower bounds of MI:

$$I(X; Y) \geq -\mathbb{E}_{p(x)}[\log p(x)] + \mathbb{E}_{p(x,y)}[\log q(x | y)].$$

The lower bound holds since the KL divergence is non-negative. For the upper bound of MI, a recent work called Contrastive Log-ratio Upper Bound (CLUB) [Cheng *et al.*, 2020] provides the following form:

$$I(X; Y) \leq \mathbb{E}_{p(x,y)}[\log q(x | y)] - \mathbb{E}_{p(x)p(y)}[\log q(x | y)],$$

where  $q(x | y)$  is approximated, e.g., by maximum likelihood estimation (MLE), of actual conditional distribution  $p(x | y)$  when  $p$  is not known.

### 3 Methodology

In this section, we propose our framework for fast adaptation, named Policy Adaptation with Decoupled Representation (PAnDR). We follow the paradigm of [Raileanu *et al.*, 2020] and consider the offline-training-online-adaptation setting with a set of training environments  $\Omega_{\text{train}} = \{\mathcal{M}_i\}_{i=1}^M$  and a set of testing environments  $\Omega_{\text{test}} = \{\mathcal{M}_j\}_{j=1}^N$ . Given a set of training policies  $\Pi_{\text{train}} = \{\pi_k\}_{k=1}^O$ , an offline experience buffer is established as  $\mathcal{D} = \bigcup_{i=1, k=1}^{M, O} D_{i,k}$ , where  $D_{i,k}$  contains the  $s, a, r$ -trajectories generated by  $\pi_k$  in  $\mathcal{M}_i$ .  $M, N, O$  denote the sizes of corresponding set. In this paper, we call a set of  $(s, a, r)$  generated in an environment as the *context* of it, and a set of  $(s, a)$  collected by a policy as its *behavior*, denoted by  $c$  and  $b$  respectively.

With offline experience buffer  $\mathcal{D}$ , PAnDR learns the representations of environment and policy through context contrast and policy recovery, respectively (Sec. 3.1). To eliminate the redundant information entangled in the two representations while retaining necessary information, we further introduce a novel approach by mutual information minimization and maximization to refine the representations (Sec. 3.2). Thereafter, we train a nonlinear neural network to approximate PDVF conditioning on both environment and policy representations and introduce the online adaptation with gradient-based policy optimization (Sec. 3.3).

#### 3.1 Self-supervised Representation Learning

The first part of PAnDR is to learn effective representations for environment and policy from offline experiences, replying on which the environment generalization and policy adaptation are carried out. We learn these representations in a self-supervised manner via context contrast and policy recovery, to capture the pivotal information of environment variation and policy behavior. We introduce the details below.

**Learning Environment Representation via Context Contrast.** By following the general idea of self-supervised Contrastive Learning [He *et al.*, 2020], we make use of the contrast of contexts collected in different training environments to extract context embedding as environment representation. For a training environment  $\mathcal{M}_i$ , we use the contexts sampled from  $D_{i,\cdot} = \bigcup_{k=1}^O D_{i,k}$  as anchor sample  $c_i$  and positive sample  $c_i^+$ ; in turn, the context sampled from the offline experiences generated in other environments, i.e.,  $D_{i'}$ , for any  $i' \neq i$ , is taken as negative sample  $c_i^-$ .

With the contrastive context samples described above, we train a context encoder  $\phi_e$  that maps context to embedding

$z_e$  by ensuring the representation of similar sample pairs (anchors and their positives) stay close to each other while dissimilar ones (anchors and their negatives) are far apart. Formally, the context encoder is optimized with InfoNCE loss similar to [He *et al.*, 2020; Laskin *et al.*, 2020]:

$$\mathcal{L}_{\text{CC}}(\phi_e, W) = -\mathbb{E}_{c, c^+, \{c^-\} \sim \mathcal{D}} \left[ z_e^\top W z_e^+ - \log \left( \exp(z_e^\top W z_e^+) + \sum_{\{c^-\}} \exp(z_e^\top W z_e^-) \right) \right], \quad (1)$$

where  $W$  denotes the parameter matrix updated along with  $\phi_e$ . Note that  $z_e$  is differentiable with respect to  $\phi_e$ ; while  $z_e^+, z_e^-$  are the embeddings for positive and negative context samples calculated by a momentum averaged version of  $\phi_e$ .

Compared to learning environment representation by transition-level dynamics prediction as adopted in [Raileanu *et al.*, 2020], PAnDR captures environment variations in an integral manner and emphasizes inter-environment differences by context-based contrastive learning. Thus, the environment representation can be more effective and robust, as similarly demonstrated in [Fu *et al.*, 2021].

**Learning Policy Representation via Policy Recovery.** For a training policy  $\pi_k$ , we learn the embedding of the behavior data sampled from  $D_{\cdot,k} = \bigcup_{i=1}^M D_{i,k}$  as policy representation. As shown in the left of Fig. 2, a behavior  $b$  is fed in the behavior encoder  $\phi_\pi$ , producing the embedding  $z_\pi$ . With the purpose of making  $z_\pi$  contain the essential policy information, we utilize a  $z_\pi$ -conditioned policy decoder  $\phi_\pi^{\text{dec}}$  which takes as input a state  $s$  and the embedding  $z_\pi$ , then deterministically outputs a prediction of the actual action made by  $\pi$ . We call this *policy recovery* (or Behavioral Cloning [Grover *et al.*, 2018]). With the offline experiences,  $\phi_\pi, \phi_\pi^{\text{dec}}$  are updated by minimizing the  $l_2$  prediction loss:

$$\mathcal{L}_{\text{PR}}(\phi_\pi, \phi_\pi^{\text{dec}}) = \mathbb{E}_{s,a,b \sim \mathcal{D}} [\|a - \phi_\pi^{\text{dec}}(s, \phi_\pi(b))\|_2^2]. \quad (2)$$

#### 3.2 MI-based Representation Refinement

A consequent concern on the representations obtained above is the redundant and entangled information in them. This is because the representations for both environment and policy are learned from the offline experiences. The entangled information exists in both representations will negatively affect the generalization among environments and policies, thus hinders value approximation and policy adaptation in turn.

To address this problem, we propose a novel information-theoretic approach to refine the learned representations. The key idea is 1) to eliminate the redundant information contained in environment and policy representations for the decoupleness, while 2) to ensure the essential information in their joint representation are completely retained in the two representations. This is achieved by two ways of mutual information (MI) optimization, as illustrated in the middle part of Fig. 2.

**MI Minimization for Decoupleness.** Firstly, for the decoupleness, we are interested in reducing the entangled information in the representations of environment and policy. To be specific, we minimize the MI between the embeddings  $z_e$  and  $z_\pi$  that are obtained from context and behavior data sampled from the offline experiences, i.e.,  $\min I(z_e; z_\pi)$ . We then

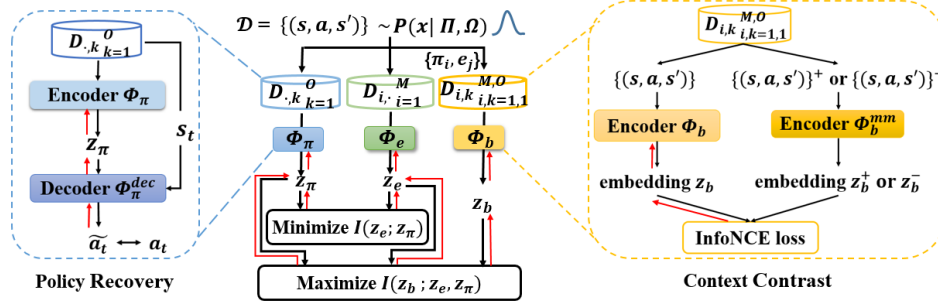


Figure 2: Illustration of learning environment representation  $z_e$  and policy representation  $z_\pi$  from the offline experiences  $\mathcal{D}$  collected by training environments  $\Omega$  and policies  $\Pi$ . With different views of  $\mathcal{D}$ , i.e.,  $D_{\cdot,k}$ ,  $D_{i,\cdot}$  and  $D_{i,k}$ , the representations for environment, policy and their joint ( $z_b$ ) are obtained by encoders  $\phi_e$ ,  $\phi_\pi$ ,  $\phi_b$  respectively, which are learned by policy recovery (left) and context contrast (right). Thereafter, the representations of environment and policy are refined through mutual information based optimization, i.e.,  $\min I(z_e; z_\pi)$  for representation decoupleness and  $\max I(z_b; z_e, z_\pi)$  for representation completeness. Gradient traces of model parameters are plotted in red.

convert the MI minimization into minimizing the CLUB upper bound of MI, as introduced in Sec. 2.2. With the approximate conditional distribution  $q_{\psi_1}$  parameterized by  $\psi_1$ , our loss function for representation decoupleness (RD) is formulated as:

$$\mathcal{L}_{\text{RD}}(\phi_e, \phi_\pi) = \mathbb{E}_{p(z_e, z_\pi)} [\log q_{\psi_1}(z_e | z_\pi)] - \mathbb{E}_{p(z_e)p(z_\pi)} [\log q_{\psi_1}(z_e | z_\pi)]. \quad (3)$$

Concretely, as to any  $D_{i,k}$ , we use the embedding pairs of  $c_i, b_k \sim D_{i,k}$  for the samples from joint distribution  $p(z_e, z_\pi)$ ; and use those of  $c_i \sim D_{i,\cdot}, b_k \sim D_{\cdot,k}$  for the marginal distributions  $p(z_e), p(z_\pi)$ . Note that  $\psi_1$  is not updated by  $\mathcal{L}_{\text{RD}}$ , while the approximate conditional distribution  $q_{\psi_1}$  is trained through MLE by minimizing the following loss:

$$\mathcal{L}_{\text{MLE}}(\psi_1) = -\mathbb{E}_{p(z_e, z_\pi)} [\log q_{\psi_1}(z_e | z_\pi)], \quad (4)$$

where the joint distribution is sampled in the same way as in Eq. 3. The updates defined in Eq. 3 and Eq. 4 are performed iteratively:  $q_{\psi_1}$  approximates the conditional distribution with respect to the embeddings given by  $\phi_e, \phi_\pi$ ; and in turn  $q_{\psi_1}$  is used in the estimation of CLUB. We refer the readers to the original paper of CLUB [Cheng *et al.*, 2020] for better understanding.

**MI Maximization for Completeness.** Secondly, we expect the information contained in the joint representation of environment and policy can be completely retained in both environment and policy representations. This is also necessary especially when the decoupleness described above is required, in case that the vital information of environment variation or policy behavior is overly reduced during MI minimization. This is achieved by maximizing the MI between the concatenation of the embeddings  $z_e, z_\pi$  and a target environment-policy joint representation  $z_b$  which contains complete information, i.e.,  $\max I(z_e, z_\pi; z_b)$ . We resort to extracting the target representation  $z_b$  also from the offline experiences. Similar but distinct to the environment representation,  $z_b$  is learned by context-behavior joint contrast, i.e., only the context-behavior data sampled from the same  $D_{i,k}$  are regarded as mutually positive samples. We omit the similar formulations here and an illustration is shown in the right part of Fig. 2. Thereafter, we maximize the variational lower bound

of  $I(z_e, z_\pi; z_b)$ , i.e., minimize the loss function for representation completeness (RC) below:

$$\mathcal{L}_{\text{RC}}(\phi_e, \phi_\pi, \psi_2) = -\mathbb{E}_{p(z_e, z_\pi, z_b)} [\log q_{\psi_2}(z_b | z_e, z_\pi)], \quad (5)$$

where  $q_{\psi_2}$  is the variational distribution parameterized by  $\psi_2$ , and the joint distribution are sampled by  $c_i, b_k \sim D_{i,k}$ . Note that the negative entropy term  $-\mathcal{H}(z_b) = \mathbb{E}_{p(z_b)} [\log p(z_b)]$  irrelevant to the optimization of  $z_e, z_\pi$  is neglected in Eq. 5.

Therefore, combining the self-supervised learning losses and the MI-based refinement losses, the total loss function for the representations of environment and policy is:

$$\mathcal{L}_{\text{Total}}(\phi_e, \phi_\pi) = \mathcal{L}_{\text{CC}} + \mathcal{L}_{\text{PR}} + \alpha \mathcal{L}_{\text{RD}} + \beta \mathcal{L}_{\text{RC}}, \quad (6)$$

where  $\alpha, \beta$  are the balancing weights. The pseudocode of overall representation training is in Alg. 1 in Appendix.

### 3.3 Value Approximation and Policy Adaptation

So far, we have introduced the training process of environment and policy representations on the offline experiences. With the trained representations, we are ready to introduce the approximation of PDVF [Raileanu *et al.*, 2020] and the policy optimization during online adaptation.

**Value Function Network Approximation.** As described in Section 1, to evaluate the values of different environment-policy configurations, PDVF additionally conditions on the representations of environment and policy. Distinct from the quadratic approximation adopted in [Raileanu *et al.*, 2020], PANDR leverages a typical multi-layer nonlinear neural network, denoted by  $\mathbb{V}_\theta(s, z_e, z_\pi)$  with parameter  $\theta$ . This setting frees the function expressivity constrained by quadratic form. We suggest that this is important because the value approximation for multiple environments and policies is much more difficult than the conventional situation (e.g.,  $V^\pi$ ). Achieving better value approximation is crucial to the following policy adaptation.

For any offline experience  $D_{i,k}$ , we can obtain the corresponding representations  $z_e, z_\pi$ , conditioning on which  $\mathbb{V}_\theta$  is trained by Monte Carlo method [Sutton and Barto, 1998]:

$$\mathcal{L}_V(\theta) = \mathbb{E}_{s, G(s), c, b \sim \mathcal{D}} [(G(s) - \mathbb{V}_\theta(s, z_e, z_\pi))^2], \quad (7)$$

where  $G(s) = \sum_{t=0}^T \gamma^t r_t$  letting  $s_0 = s$  is the Monte Carlo return. The pseudocode is provided in Alg. 2 in Appendix.




Figure 3: The x-axis represents five test environments, and the y-axis is average return. Adaptation performance of different algorithms in four domains. Experimental results of algorithm generalization in new environment. Our results are even better than PPO in some environments, which proves that our method has a better generalization ability by learning in the training environment. For the PDVF algorithm, we use the source code provided by the authors.

Domain	Train/Test Env. No.	Variation
Spaceship-Charge	15 / 5	The strength of electric field
Ant-Wind	15 / 5	Direction of wind
Swimmer-Fluid	15 / 5	Direction of the fluid
HalfCheetah-Mass	8 / 5	Joint damping coefficient

Table 1: Domains in our experiments for fast adaptation.

### Policy Adaptation via Representation Gradient Ascent.

After training  $\nabla_{\theta}$  network well, in principle we can evaluate the values for any environment-policy configuration to a certain degree. For online adaptation, we first sample a training policy from  $\Pi_{\text{train}}$  and use it to interact with the testing environment for a few steps. Then the embedding of testing environment  $z_e^{\text{on}}$  can be inferred with the collected context. Next, starting from an arbitrary policy representation  $z_{\pi}^0$  (randomly initialized by default), by gradient ascent (GA), we can optimize the policy representation along the direction of increasing values of  $\nabla_{\theta}$  with learning rate  $\eta$ :

$$z_{\pi}^{t+1} = z_{\pi}^t + \eta \nabla_{z_{\pi}^t} \nabla_{\theta}(s, z_e^{\text{on}}, z_{\pi}^t). \quad (8)$$

After a certain number of GA, the obtained policy representation can be used to interact with the testing environment with the help of the  $z_{\pi}$ -conditioned policy decoder  $\phi_{\pi}^{\text{dec}}$  (described in Sec. 3.1), which is trained offline to predict policy’s action. We adopt the best-performing policy representation as  $z_{\pi}^*$  encountered during the GA process for our adapted policy in our experiments.

It is worth noting that the GA optimization itself is zero-shot, i.e., without the overhead of interacting with the testing environment. The pseudocode is in Alg. 3 in Appendix.

## 4 Experiment

In our experiments, we mainly aim at investigating the following questions: **(Q1)** Whether PAnDR outperforms existing algorithms for fast adaptation? **(Q2)** How do contrastive environment representation and MI-based refinement contribute to the performance of PAnDR? **(Q3)** Whether nonlinear neural approximation of PDVF and gradient-based policy adaptation improve the quadratic counterpart?

### 4.1 Setup

We conduct experiments in four continuous control domains. All the experimental details can be found in Appendix B.

**Domains and Variations.** Table 1 shows the fast adaptation domains used in our experiments, where the first 3 domains are from [Raileanu *et al.*, 2020] and the last one is

from [Lee *et al.*, 2020]. In each domain, we have a variety of environments controlled by corresponding variations. These variations determine the dynamics and reward functions of environment. Therefore, fast adaptation is expected to be achieved against these changes. The number of training and testing environments are provided in Table 1.

**Offline Experiences and Adaptation Budget.** We follow the same data setting as [Raileanu *et al.*, 2020]. For offline experiences, a policy is optimized by PPO [Schulman *et al.*, 2017] in each training environments. The obtained policies are training policies that serve as collectors. In general, such collector policies can be obtained from other ways. For each domain, 50-episode interaction trajectories are collected by each combination of training environment and policy, e.g.,  $15 * 15 * 50$  episodes of experience for Ant-Wind. For online adaptation, 50 steps of interaction are collected in the testing environment by one of the training policies. For PAnDR, 100 steps of GA is used by default.

**Evaluation Criteria.** The main evaluation criterion is *adaptation performance*, calculated as the undiscounted cumulative rewards of the policy obtained after online adaptation. Moreover, for PAnDR, we analyze the evolvement of adaptation performance against the number of GA adaptation steps (in Sec. 4.3). The results reported in our experiments are means and error bars of a standard deviation calculated with 7 independent trials.

### 4.2 Performance Comparison

In this section, we compare the adaptation performance between PAnDR and representative existing algorithms to answer **Q1** within the domains introduced above.

**Baselines.** We consider the following baselines. 1) PPO: We use the PPO policy which is trained online with  $3e6$  steps in testing environments as a *powerful* baseline. 2) MAML [Finn *et al.*, 2017]: We use MAML as a representative of gradient-based Meta RL. Note MAML (as well as most Meta RL algorithms) performs online meta training. This means in principle it is *not applicable* to the offline training experiences. 3) PDVF [Raileanu *et al.*, 2020]: Since we propose PAnDR based on the same paradigm of PDVF, we use PDVF as a closely related baseline. Moreover, to our knowledge, PDVF is the state-of-the-art fast adaptation algorithm in offline-training-online-adaptation setting. Note that in this section, we use PDVF to denote the method name of [Raileanu *et al.*, 2020] rather than the extended value function. For all above baselines, we use the official source codes




Figure 4: Ablations for MI-based representation refinement.

provided by PDVF paper [Raileanu *et al.*, 2020]. For PAnDR, we also implement upon these source codes with almost the same configuration of hyperparameters and structures for a valid comparison.

**(A1) Results.** As shown in Fig. 3, for the simple Spaceship domain, the overall performance of PAnDR is significantly better than PDVF and slightly outperforms MAML. Further in the three more difficult MuJoCo domains, the leading edge of PAnDR is much more evident and consistent compared with the two baseline methods, while MAML loses its competitiveness. We consider that this is because the preferable representations obtained via self-supervised learning and the MI-based refinement are more effective and robust than those of PDVF. Moreover, with less-constrained approximation, the more accurate value estimation achieved improves the optimality of policy obtained during adaptation.

It is worth noting that, perhaps somewhat surprisingly, PAnDR shows competitive and even better performance when comparing with the powerful PPO baseline trained in the testing environments. Instead of requiring millions of interactions to optimize a policy in the testing environments, PAnDR only needs limited interacting experiences for fast adaptation. This further reveals the superiority and potential of PAnDR.

### 4.3 Ablation Study and Analysis

In this section, we give the answers to Q2 and Q3. We provide main results below, and the complete results along with detailed analysis are given in Appendix A.

**(A2-I) MI-based Representation Refinement.** To verify the contribution of the MI minimizing and maximizing in representation refinement of PAnDR, we compare PAnDR with its several variants including: 1) PAnDR without MAX MI: removes MI maximizing refinement; 2) PAnDR without MIN MI: removes MI minimizing refinement; 3) PAnDR without MI: removes both MI maximizing and minimizing refinements. As shown in Fig.4, both MI minimizing and maximizing positively contribute to the improvement of PAnDR’s overall performance. It demonstrates that both MI-based refinement components are indispensable in PAnDR.

**(A2-II) Context Contrast for Environment Representation.** Fig.5, presents the efficacy of context contrast in comparison with dynamics prediction upon both types of  $\nabla_{\theta}$  in PDVF and PAnDR. The associated variants are shown in Table 2. As in Fig.5, PDVF with CER has superior performance than original PDVF while PAnDR without MI outperforms PAnDR with NA-GA. Thus, context contrast is more favorable in both cases of  $\nabla_{\theta}$  compared with dynamics prediction.

**(A3-I)  $\nabla_{\theta}$  Approximation & GA Adaptation.** With the variants in Table 2, we compares: PAnDR v.s. PAnDR with QA, and PDVF v.s. PDVF with NA-GA. The results in Fig.5 verify that there is a clear advantage of using NN nonlinear




Figure 5: Efficacy analysis for contrast environment representation (CER) and neural approximation of PDVF with GA adaptation (NA-GA) in contrast to quadratic approximation (QA).

Alg. / Variant	Env. Repr.	$\nabla_{\theta}$ Approx. & Policy Opt.	If Use MI
PDVF	DP	Quadratic & Close-form	✗
PDVF with CER	CL	Quadratic & Close-form	✗
PDVF with NA-GA	DP	NN Nonlinear & GA	✗
PAnDR without MI	CL	NN Nonlinear & GA	✗
PAnDR with QA	CL	Quadratic & Close-form	✓
PAnDR	CL	NN Nonlinear & GA	✓

Table 2: Comparison of the variants with respect to three algorithmic factors (Fig. 5). Abbreviations are used, i.e., Dynamics Prediction (DP), Contrastive Learning (CL).

approximation of  $\nabla_{\theta}$  along with GA optimization over the quadratic counterpart.

**(A3-II) Adaptation Performance against GA Steps.** Since there is no longer a closed-form solution to optimization available for  $\nabla_{\theta}$  with NN nonlinear approximation, we investigate how consecutive steps of GA affect the adaptation performance of PAnDR. The results for various choices of GA step number are provided in Table 3-8 in Appendix A.3. Moreover, we provide detailed analysis for useful insights.

## 5 Conclusion and Discussion

This paper focus on fast adaptation problem with only offline experiences on training environments available. Following the paradigm of [Raileanu *et al.*, 2020], we propose a new algorithm, PAnDR, that consists of effective environment and policy representations and better policy optimization during online adaptation. In essence, policy adaptation in PAnDR (and PDVF) is a process of policy representation optimization and decoding. Both the two parts are not well studied at present. Optimization in policy representation space is non-trivial as there is a lack of knowledge on the landscape of the space; while the decoding of such latent representation can be invalid without carefully handling the uncertainty [Notin *et al.*, 2021]. We consider that they are significant and there is much room for improvement. In addition, advances in environment and policy representation learning are expected to further improve the performance of PAnDR.

## Acknowledgments

This work is supported by the National Science Fund for Distinguished Young Scholarship of China (No. 62025602), the National Key R&D Program of China (Nos. 2018AAA0100900), the National Natural Science Foundation of China (Nos. U1803263, 62106172 and 11931015), the Key Technology Research and Development Program of Science and Technology-Scientific and Technological Innovation Team of Shaanxi Province (Grant No. 2020TD-013), the XPLORER PRIZE, the New Generation of Artificial Intelligence Science and Technology Major Project of Tianjin (Grant No.: 19ZXZNGX00010), and the Science and

Technology on Information Systems Engineering Laboratory (Grant No. WDZC20205250407).

## References

- [Alemi *et al.*, 2017] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy. Deep variational information bottleneck. In *ICLR*, 2017.
- [Andrychowicz *et al.*, 2020] M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. W. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, and W. Zaremba. Learning dexterous in-hand manipulation. *IJRR*, 39(1):3–20, 2020.
- [Badia *et al.*, 2020] A. P. Badia, B. Piot, S. Kapturowski, P. Sprechmann, A. Vitvitskyi, Z. Daniel Guo, and C. Blundell. Agent57: Outperforming the atari human benchmark. In *ICML*, volume 119, pages 507–517, 2020.
- [Cheng *et al.*, 2020] P. Cheng, W. Hao, S. Dai, J. Liu, Z. Gan, and L. Carin. Club: A contrastive log-ratio upper bound of mutual information. In *ICML*, pages 1779–1788, 2020.
- [Duan *et al.*, 2016] Y. Duan, J. Schulman, X. Chen, P. Bartlett, I. Sutskever, and P. Abbeel. RL<sup>2</sup>: Fast reinforcement learning via slow reinforcement learning. *CoRR*, abs/1611.02779, 2016.
- [Faccio *et al.*, 2021] F. Faccio, L. Kirsch, and J. Schmidhuber. Parameter-based value functions. In *ICLR*, 2021.
- [Fakoor *et al.*, 2020] R. Fakoor, P. Chaudhari, S. Soatto, and A. J. Smola. Meta-q-learning. In *ICLR*, 2020.
- [Finn *et al.*, 2017] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [Fu *et al.*, 2021] H. Fu, H. Tang, J. Hao, C. Chen, X. Feng, D. Li, and W. Liu. Towards effective context for meta-reinforcement learning: an approach based on contrastive learning. In *AAAI*, pages 7457–7465, 2021.
- [Grover *et al.*, 2018] A. Grover, M. Al-Shedivat, J. K. Gupta, Y. Burda, and H. Edwards. Learning policy representations in multiagent systems. In *ICML*, volume 80, pages 1797–1806, 2018.
- [Guo *et al.*, 2020] Z. Guo, B. Pires, B. Piot, J. Grill, F. Althché, R. Munos, and M. Azar. Bootstrap latent-predictive representations for multitask reinforcement learning. In *ICML*, volume 119, pages 3875–3886, 2020.
- [Harb *et al.*, 2020] J. Harb, T. Schauland D. Precup, and P. Bacon. Policy evaluation networks. *CoRR*, abs/2002.11833, 2020.
- [He *et al.*, 2020] K. He, H. Fan, Y. Wu, S. Xie, and R. B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735, 2020.
- [Kirk *et al.*, 2021] R. Kirk, A. Zhang, E. Grefenstette, and T. Rocktäschel. A survey of generalisation in deep reinforcement learning. *CoRR*, abs/2111.09794, 2021.
- [Laskin *et al.*, 2020] M. Laskin, A. Srinivas, and P. Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. In *ICML*, pages 5639–5650, 2020.
- [Lee *et al.*, 2020] K. Lee, Y. Seo, S. Lee, H. Lee, and J. Shin. Context-aware dynamics model for generalization in model-based reinforcement learning. In *ICML*, pages 5757–5766, 2020.
- [Ma *et al.*, 2021] Y. Ma, X. Hao, J. Hao, J. Lu, X. Liu, X. Tong, M. Yuan, Z. Li, J. Tang, and Z. Meng. A hierarchical reinforcement learning based optimization framework for large-scale dynamic pickup and delivery problems. In *NeurIPS*, pages 23609–23620, 2021.
- [Notin *et al.*, 2021] P. Notin, J. M. Hernández-Lobato, and Y. Gal. Improving black-box optimization in VAE latent space using decoder uncertainty. *CoRR*, abs/2107.00096, 2021.
- [Packer *et al.*, 2018] C. Packer, K. Gao, J. Kos, P. Krähenbühl, V. Koltun, and D. Song. Assessing generalization in deep reinforcement learning. *CoRR*, abs/1810.12282, 2018.
- [Peng *et al.*, 2018] X. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *ICRA*, pages 1–8, 2018.
- [Poole *et al.*, 2019] B. Poole, S. Ozair, A. Oord, A. Alemi, and G. Tucker. On variational bounds of mutual information. In *ICML*, volume 97, pages 5171–5180, 2019.
- [Raileanu *et al.*, 2020] R. Raileanu, M. Goldstein, A. Szlam, and R. Fergus. Fast adaptation to new environments via policy-dynamics value functions. In *ICML*, pages 7920–7931, 2020.
- [Rakelly *et al.*, 2019] K. Rakelly, A. Zhou, C. Finn, S. Levine, and D. Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *ICML*, volume 97, pages 5331–5340, 2019.
- [Schulman *et al.*, 2017] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Sutton and Barto, 1998] R. S. Sutton and A. G. Barto. *Reinforcement learning - an introduction*. Adaptive computation and machine learning. MIT Press, 1998.
- [Tang *et al.*, 2020] H. Tang, Z. Meng, J. Hao, C. Chen, D. Graves, D. Li, C. Yu, H. Mao, W. Liu, Y. Yang, and L. Wang. What about inputting policy in value function: Policy representation and policy-extended value function approximator. *CoRR*, abs/2010.09536, 2020.
- [Yarats *et al.*, 2021] D. Yarats, I. Kostrikov, and R. Fergus. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *ICLR*, 2021.
- [Zheng *et al.*, 2019] Y. Zheng, C. Fan, X. Xie, T. Su, L. Ma, J. Hao, Z. Meng, Y. Liu, R. Shen, and Y. Chen. Wuji: Automatic online combat game testing using evolutionary deep reinforcement learning. In *ASE*, pages 772–784, 2019.
- [Zheng *et al.*, 2021] Y. Zheng, Y. Liu, X. Xie, Y. Liu, L. Ma, J. Hao, and Y. Liu. Automatic web testing using curiosity-driven reinforcement learning. In *ICSE*, pages 423–435, 2021.