

# In-Context Policy Adaptation via Cross-Domain Skill Diffusion

Minjong Yoo, Woo Kyung Kim, Honguk Woo\*

Department of Computer Science and Engineering, Sungkyunkwan University  
 {mjyoo2, kwk2696, hwoo}@skku.edu

## Abstract

In this work, we present an in-context policy adaptation (ICPAD) framework designed for long-horizon multi-task environments, exploring diffusion-based skill learning techniques in cross-domain settings. The framework enables rapid adaptation of skill-based reinforcement learning policies to diverse target domains, especially under stringent constraints on no model updates and only limited target domain data. Specifically, the framework employs a cross-domain skill diffusion scheme, where domain-agnostic prototype skills and a domain-grounded skill adapter are learned jointly and effectively from an offline dataset through cross-domain consistent diffusion processes. The prototype skills act as primitives for common behavior representations of long-horizon policies, serving as a lingua franca to bridge different domains. Furthermore, to enhance the in-context adaptation performance, we develop a dynamic domain prompting scheme that guides the diffusion-based skill adapter toward better alignment with the target domain. Through experiments with robotic manipulation in Metaworld and autonomous driving in CARLA, we show that our ICPAD framework achieves superior policy adaptation performance under limited target domain data conditions for various cross-domain configurations including differences in environment dynamics, agent embodiment, and task horizon.

## Introduction

Reinforcement learning (RL) has exhibited remarkable success in addressing sequential decision-making problems. However, a fundamental challenge emerges in the context of policy adaptation, the process of applying a policy trained in a source domain via RL to a different target domain, particularly for complex, long-horizon tasks (Zhu et al. 2023; Zhao, Queralta, and Westerlund 2020; Hua et al. 2021; He et al. 2024; Hakhamaneshi et al. 2022). This challenge escalates under practical restrictions that data availability is limited and direct interaction with the target domain is prohibited.

To address this challenge, we explore skill-based RL approaches (Pertsch, Lee, and Lim 2021; Nam et al. 2022) in cross-domain environments, with the aim of facilitating the in-context adaptation of skill-based policies to different target domains under limited target data. We present an in-context policy adaptation framework, namely ICPAD. As

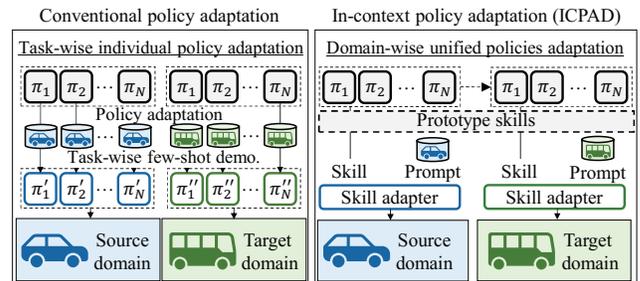


Figure 1: A middle-tier strategy for cross-domain policy adaptation in ICPAD: as shown on the right, in ICPAD, policies  $\pi$  learned in the source domain with domain-agnostic prototype skills (in the middle-tier) are capable of adapting to the target domain in an in-context unified way by exploiting few-shot target domain data. Unlike the conventional policy adaptation relying on *policy-wise* model updates (shown on the left), ICPAD allows for simultaneous adaptation of multiple policies to the target domain through *domain-wise* middle-tier adaptation.

illustrated in Figure 1, ICPAD facilitates policy adaptation to target domains in a unified way, by learning transferable prototype skills and skill adapters offline, and implementing a middle-tier adaptation strategy. Unlike conventional policy adaptation approaches, which typically require task-specific individual policy transfers, ICPAD offers a notable advantage in cross-domain multi-task environments. Its effectiveness stems from prototype skills that function as the middle-tier adaptation layer, linking policies and domains. These skills, which support intermediate representations or functional primitives of policies, facilitate their translation into domain-specific actions through dedicated skill adapters. In doing so, we adapt skill-based RL approaches to the context of cross-domain few-shot policy adaptation. That is, source domains have abundant data, while target domains are restricted to minimal data.

In the offline learning phase of ICPAD, we first establish the middle-tier adaptation layer, prototype skills, which encapsulate expert behavior patterns and are transferable across domains. Specifically, we develop a cross-domain skill diffusion scheme, incorporating the distribution matching-based cross-domain consistency into the skill diffusion process.

\*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This scheme allows for the domain alignment of action sequences, which are generated specifically for the target domain from the prototype skills. Then, using these prototype skills, which are established domain-agnostically from the offline dataset, RL policies are trained in the source domain.

In the in-context adaptation phase of ICPAD, where the policies can be rapidly adapted to the target domain without model updates, we use a dynamic domain prompting scheme. It provides robust online guidance on the skill diffusion process toward enhanced alignment with the target domain, using only a few target data for in-context adaptation. These learning phases of ICPAD enable the in-context adaptation capabilities of skill-based policies.

Through a series of experiments encompassing robotic manipulation in Metaworld (Yu et al. 2020) and autonomous driving in CARLA (Dosovitskiy et al. 2017), we demonstrate that ICPAD is applicable for a variety of cross-domain configurations including varied environment dynamics and different embodiment conditions. Furthermore, the policies learned via the framework (ICPAD policies) consistently outperform other state-of-the-art skill-based meta-RL methods. For example, in autonomous driving scenarios, ICPAD policies exhibit superior performance, achieving 11.6% to 21.6% higher normalized returns than the most competitive baseline DCMRL (He et al. 2024).

Our main contributions are summarized as follows.

- We present a novel in-context policy adaptation ICPAD framework to facilitate domain-wise rapid adaptation of skill-based policies across diverse domains without the need for online model updates.
- We develop the cross-domain skill diffusion scheme, designed to integrate cross-domain consistency learning and skill diffusion, thereby achieving domain-agnostic prototype skills robustly and enabling their effective translation to domain-specific actions.
- We adopt the dynamic domain prompting scheme, designed to enhance the adaptability of skill-based policies to the target domain.
- Through extensive experiments with Metaworld and CARLA, we demonstrate ICPAD’s effectiveness in a variety of cross-domain settings, establishing its superiority over several state-of-the-art methods.

## Preliminary

**Skill-based RL.** With the goal of accelerating policy learning in complex environments, skill-based RL approaches leverage offline imitation learning on task-agnostic expert datasets, establishing a set of shared skills that capture expert behavior patterns. As implemented in (Pertsch, Lee, and Lim 2021), these approaches encompass several models; a skill encoder maps expert sub-trajectories to a skill latent space and a skill decoder converts skill embeddings to action sequences. Unlike existing skill-based RL approaches that require online policy learning within target domains, our focus is on in-context policy adaptation to given target domains, prohibiting online environment interactions.

**Diffusion model.** The diffusion model (HO, Jain, and Abbeel 2020) frames the generation process as an iterative denoising

procedure, where a noisy input  $x^K$  sampled from a standard normal distribution  $p(x^K)$  is gradually denoised to recover the noiseless data  $x^0$ . That is

$$p(x^0) = p(x^K) \prod_{k=1}^K p(x^{k-1}|x^k), \quad (1)$$

where  $k$  is the denoising timestep. As training a diffusion model  $\epsilon$  involves optimizing a variational lower bound of  $p(x^0)$ , the below surrogate loss can be used.

$$\mathbb{E}_{k \sim [1, K], \eta \sim \mathcal{N}(0, I)} [\|\eta - \epsilon(x^k, k)\|_2^2] \quad (2)$$

The noisy input  $x^k$  is generated by adding the Gaussian noise to the original data  $x^0$  as

$$x^k = \sqrt{\bar{\alpha}^k} x^0 + \sqrt{1 - \bar{\alpha}^k} \eta \quad (3)$$

where  $\bar{\alpha}^k = \prod_{s=1}^k \alpha^s$  is a variance schedule. This process can be viewed as estimating the noise  $\eta$  that is added to the original data  $x^0$  to generate the noisy input  $x^k$ . For generation, the diffusion model uses the iterative denoising chain such as

$$x^{k-1} = \frac{1}{\sqrt{\alpha^k}} \left( x^k - \frac{1 - \alpha^k}{\sqrt{1 - \bar{\alpha}^k}} \epsilon(x^k, k) \right). \quad (4)$$

## Problem Formulation

### Cross-domain policy adaptation without model updates.

We consider the challenge of efficiently adapting RL policies to diverse target domains in long-horizon multi-task environment with limited target data. Specifically, we explore skill-based RL approaches utilizing a task-agnostic expert dataset collected from multiple domains. The dataset enables the creation of shareable, pre-trained skills, where each skill captures a reusable action sequence or behavior pattern.

Our goal is to facilitate in-context adaptation of policies learned on these pre-trained skills, allowing for rapid policy deployment in the target domain without the need for gradient-based model updates or intensive environment interactions. This adaptation is conducted under limited target data conditions, leveraging only few-shot target domain data (i.e., no more than 5 trajectories across multiple tasks). The data requirements for offline skill learning are consistent with those of existing skill-based RL (Pertsch, Lee, and Lim 2021; Shi, Lim, and Lee 2022). For the target domain, the data requirements align with the few-shot adaptation research (Xu et al. 2022; Hakhmaneshi et al. 2022; He et al. 2024).

## Our Approach

### Overall Framework

To achieve in-context policy adaptation across domains, we present the ICPAD framework involving two phases: (i) offline learning for cross-domain skill diffusion and (ii) in-context adaptation by dynamic domain prompting. (i) In the offline learning phase, we first establish domain-agnostic prototype skills and a domain-grounded skill adapter by leveraging an offline dataset collected from multiple domains. (ii) In the latter phase of in-context adaptation, we facilitate the adaptation of policies learned with the prototype skills, by incorporating appropriate domain information into the context

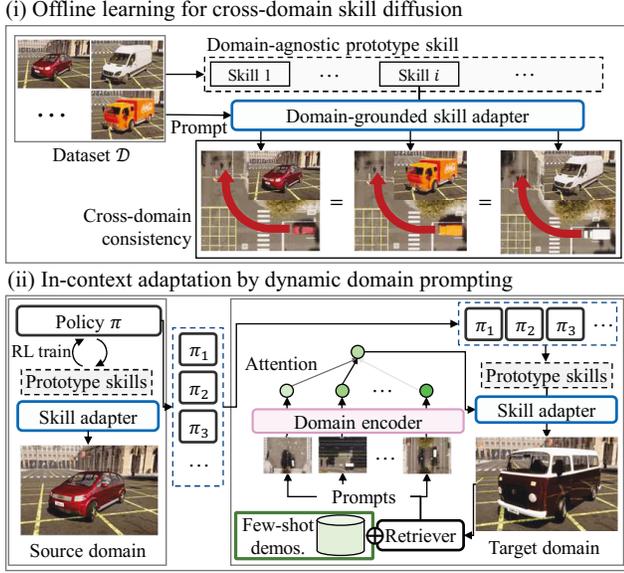


Figure 2: Offline learning for cross-domain skill diffusion and in-context adaptation by dynamic domain prompting in ICPAD: (i) In the offline learning phase, domain-agnostic prototype skills are learned jointly with a domain-grounded skill adapter that is prompted to generate consistent action sequences across domains. (ii) In the in-context adaptation phase, policies  $\pi$  learned from the source domain are adapted to unseen target domains, facilitated by the skill adapter being prompted through retrieval-based attention.

of diffusion-based skill translation, ensuring better alignment with the target domain. This dynamic domain prompting enables effective online policy adaptation, allowing for rapid deployment in diverse target domains.

Figure 2 illustrates the concept of ICPAD framework. In the offline learning, we ensure that the distribution of prototype skills remains consistent across the domains, enabling the skill adapter to generate action sequences that accurately reflect the characteristics of both the skills and domains. We also train the domain encoder through contrastive learning, which is used to dynamically retrieve proper prompts and applies attention to their embedding for in-context adaptation. Algorithm 1 lists the overall procedure of ICPAD.

## Cross-domain Skill Diffusion

**Domain-agnostic prototype skills.** To learn the prototype skills, we adapt the skill extraction approach (Nam et al. 2022) using the diffusion models with domain embedding. The skill encoder  $\Phi_E$  takes an  $H$ -step sub-trajectory  $\tau = (s, a)$  and returns the distribution of prototype skill embedding  $p(z|s, a)$  where  $z \in \mathcal{Z}$  is the skill embedding. The domain encoder  $\Psi_E$  takes an  $H$ -step sub-trajectory to return the distribution of domain embedding  $p(d|s, a)$ .

$$\Phi_E : (s, a) \mapsto p(z|s, a), \Psi_E : (s, a) \mapsto p(d|s, a) \quad (5)$$

Here,  $s$  and  $a$  denote a state sequence  $(s_t, s_{t+1}, \dots, s_{t+H})$  and an action sequence  $(a_t, a_{t+1}, \dots, a_{t+H})$ , respectively.

## Algorithm 1: Learning procedure for ICPAD framework

```

1: // Offline learning for cross-domain skill diffusion
2: Initialize skill encoder  $\Phi_E$ , adapter  $\Phi_A$ , prior  $\Phi_R$ , domain
   encoder  $\Psi_E$ , Dataset  $\mathcal{D}$ 
3: while not converge do
4:   Sample batches  $\mathcal{B}, \mathcal{B}', \mathcal{B}_C \sim \mathcal{D}$ 
5:   Calculate the loss  $\mathcal{L}_{\text{skill}}(\Phi_E, \Phi_A)$  in (7) using  $\mathcal{B}$ 
6:   Calculate the loss  $\mathcal{L}_{\text{cross-E}}(\Phi_E, \Phi_R)$  in (9) using  $\mathcal{B}$ 
7:   Calculate the loss  $\mathcal{L}_{\text{cross-A}}(\Phi_A)$  in (10) using  $\mathcal{B}'$ 
8:   Calculate the loss  $\mathcal{L}_{\text{con}}(\Psi_E)$  in (11) using  $\mathcal{B}_C$ 
9:   Update  $\Phi_E, \Phi_A, \Phi_R, \Psi_E$  based on the sum of the losses
10: end while
11: // Iteration for each task (Line 12-24)
12: // In-context adaptation by dynamic domain prompting
13: Initialize skill-based policy  $\pi$ , dynamic prompting func-
   tion  $\Psi_D$  by  $\Psi_E$ , source env.  $env$ , replay buffer  $\mathcal{D}_R$ 
14: while not converge do
15:    $s = env.reset()$ 
16:   loop
17:      $z \sim \pi(\cdot|s), d = \Psi_D(\mathcal{D}_R, s^h)$  in (13)
18:      $a = (a_0, \dots, a_H) = \Phi_A(s, d, z)$ 
19:     Execute  $a$  in  $env$  and get  $s$ , add transitions in  $\mathcal{D}_R$ 
20:     Sample batch  $\mathcal{B}_R \sim \mathcal{D}_R$ 
21:     Optimize  $\pi$  by (14) using  $\mathcal{B}_R$ 
22:   end loop
23: end while
24: In-context adaptation with  $\pi, \Psi_D, \Phi_A$  in (15) for target

```

The skill adapter  $\Phi_A$  produces an action sequence  $a$  conditioned by state  $s_t$ , domain embedding  $d$ , and prototype skill embedding  $z$ .

$$\Phi_A : (s_t, d, z) \mapsto a \quad (6)$$

For the skill adapter, we use diffusion models (HO, Jain, and Abbeel 2020), in which an action sequence  $a(=x^0)$  is generated from a random noisy input through an iterative denoising process. Accordingly, the skill adapter  $\Phi_A$  implemented with a diffusion model  $\epsilon$  takes additional inputs of a noised-corrupted action sequence  $x^k$  and denoising step  $k$ .

To achieve the skill encoder and adapter, we implement the **skill imitation loss** integrating the diffusion loss and a regularization loss for the skill space, similar to (Pertsch, Lee, and Lim 2021). For a batch  $\mathcal{B} = \{\tau\}$  sampled from  $\mathcal{D}$ , the skill imitation loss  $\mathcal{L}_{\text{skill}}(\Phi_E, \Phi_A)$  is defined as

$$\mathbb{E}_{z \sim \Phi_E, d \sim \Psi_E, \mathcal{B}} \left[ \|\eta - \epsilon(x^k, k, s_t, d, z)\|_2^2 + \beta D_{\text{KL}}(\Phi_E(\cdot|s, a), P(z)) \right] \quad (7)$$

where  $k \sim [1, K]$  is a denoising step,  $\eta \sim \mathcal{N}(0, I)$  is noise sampled from the standard normal distribution, and  $P(z)$  is the prior distribution for the prototype skill embedding. The regularization term  $D_{\text{KL}}$  denotes the Kullback–Leibler divergence, and  $\beta$  is a coefficient for regularization.

**Cross-domain consistency learning.** To maintain cross-domain consistency for the skill encoder and ensure domain-agnostic characteristics in prototype skills, we use the cross-domain skill prior  $\Phi_R$ . This model is responsible for inferring

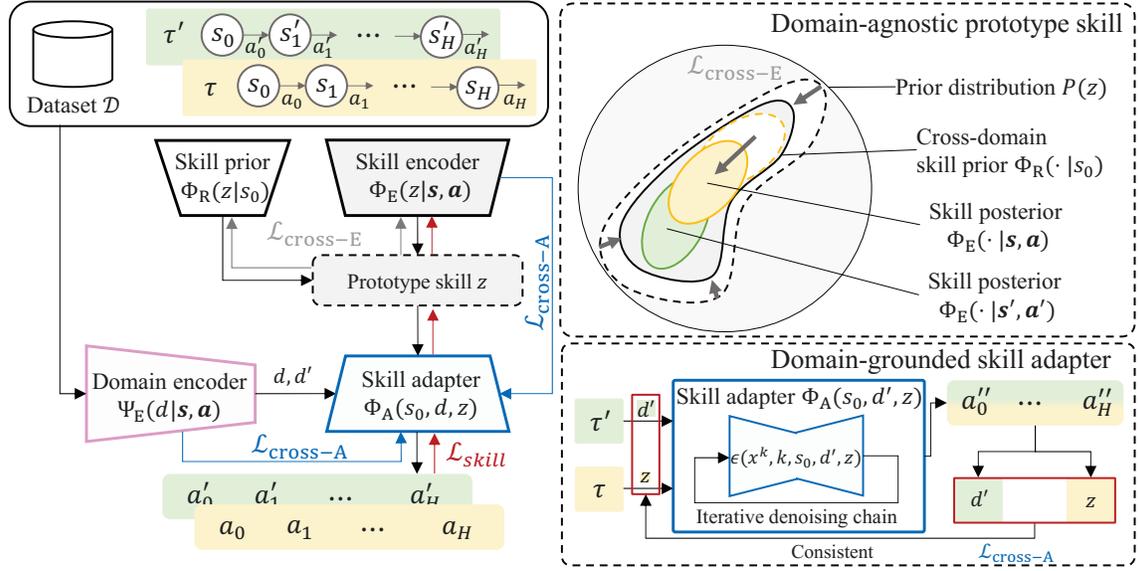


Figure 3: Cross-domain skill diffusion: we train the skill encoder  $\Phi_E$  for domain-agnostic prototype skills and the domain-grounded skill adapters  $\Phi_A$  for domain-specific skill translation across domains. These skill models are optimized using the skill imitation loss  $\mathcal{L}_{\text{skill}}$  (Eq. (7)), the cross-domain skill prior consistency loss  $\mathcal{L}_{\text{cross-E}}$  (Eq. (9)) for the skill encoder  $\Phi_E$  and skill prior  $\Phi_R$ , as well as the cross-domain action consistency loss  $\mathcal{L}_{\text{cross-A}}$  (Eq. (10)) for the skill adapter  $\Phi_A$ .

the prior distribution of skills based on states. The skill prior is formulated as

$$\Phi_R : s_t \mapsto p(z|s_t). \quad (8)$$

For a batch  $\mathcal{B}$  sampled from  $\mathcal{D}$ , the **cross-domain skill prior consistency loss**  $\mathcal{L}_{\text{cross-E}}(\Phi_E, \Phi_R)$  for the skill encoder and prior is defined as

$$\mathbb{E}_{\Phi_E, \Phi_R, \mathcal{B}} \left[ D_{\text{KL}}(\Phi_E(\cdot|s, \mathbf{a}), \text{sg}(\Phi_R(\cdot|s_t))) + \mu D_{\text{KL}}(\text{sg}(\Phi_E(\cdot|s, \mathbf{a})), \Phi_R(\cdot|s_t)) \right]. \quad (9)$$

Here,  $\text{sg}$  is a stop gradient function, and  $\mu$  is a hyperparameter. By jointly training the skill encoder and the skill prior, the skill encoder is optimized to align with the distribution of skill prior relative to states, independent of domains.

To maintain cross-domain consistency in the action sequences generated by the diffusion-based skill adapter, we align the prototype skill embedding  $z$  and domain embedding  $d$  extracted from the generated action sequences. For a batch  $\mathcal{B}' = \{\tau, \tau' = (s', \mathbf{a}')\}$  sampled from  $\mathcal{D}$ ,  $z \sim \Phi_E(s, \mathbf{a})$  and  $d' \sim \Psi_E(s', \mathbf{a}')$ , the **cross-domain action consistency loss**  $\mathcal{L}_{\text{cross-A}}(\Phi_A)$  for the skill adapter is defined as

$$\mathbb{E}_{z \sim \Phi_E, d' \sim \Psi_E, \mathcal{B}'} \left[ D_{\text{KL}}(\Psi_E(\cdot|s, \Phi_A(s_t, d', z)), \Psi_E(\cdot|s', \mathbf{a}')) + D_{\text{KL}}(\Phi_E(\cdot|s, \Phi_A(s_t, d', z)), \Phi_E(\cdot|s, \mathbf{a})) \right]. \quad (10)$$

This ensures that the action sequences generated by the skill adapter are not only appropriate for the specific domain but also consistent across different domains. Figure 3 illustrates the cross-domain skill diffusion, of which the training involves domain-agnostic prototype skills and a domain-grounded skill adapter. The joint learning approach with

the loss formulation of Eq. (7), (9), and (10) establishes the shareable prototype skill embedding across various domains, enabling the skill adapter to be effective for a wide range of domains and skills.

## Dynamic Domain Prompting

**Contrastive learning for domain encoder.** To train the domain encoder  $\Psi_E$ , we adopt contrastive learning (Schroff, Kalenichenko, and Philbin 2015). We sample a batch  $\mathcal{B}_C = \{(\tau, \tau^+), (\tau, \tau^-)\}$ , where  $(\tau, \tau^+)$  denotes a positive pair from the same domain, and  $(\tau, \tau^-)$  denotes a negative pair. Then, for positive domain embedding  $d^+ \sim \Psi_E(\cdot|\tau^+ = (s^+, \mathbf{a}^+))$  and negative domain embedding  $d^- \sim \Psi_E(\cdot|\tau^- = (s^-, \mathbf{a}^-))$ , the **contrastive loss**  $\mathcal{L}_{\text{con}}(\Psi_E)$  is defined as

$$\mathbb{E}_{d \sim \Psi_E, \mathcal{B}_C} [\max\{0, \|d - d^+\| - \|d - d^-\| + \delta\}] \quad (11)$$

where  $\delta$  is the margin of contrastive learning. Additionally, we employ the skill adapter to generate sub-trajectories from random domain embedding with the same skill embedding for negative samples.

**Retrieval-based prompt attention.** The retriever  $\Psi_R$  samples  $m$  prompts from target domain data (demonstrations)  $\mathcal{T}$ , based on a multinomial softmax distribution. The likelihood of retrieving a prompt  $\hat{\tau}_i$  is determined by the sum of Euclidean distance  $w_i$  between the state history  $s^h = (s_{t-H}, \dots, s_t)$  sampled from the policy and the states in  $\mathcal{T}$ .

$$\hat{\tau} = \{\hat{\tau}_1 = (\hat{s}_1, \hat{\mathbf{a}}_1), \dots, \hat{\tau}_m\} = \Psi_R(\mathcal{T}) \quad (12)$$

We also leverage the distance to incorporate attention into the domain embedding, ensuring that the embedding is properly chosen in the absence of data specific to the current

Domain factor	Action noise			Wind		
	Low	Medium	High	Low	Medium	High
<b>Evaluation on Metaworld</b>						
DiffBC+FT	11.3 ± 0.9%	8.2 ± 0.5%	6.0 ± 0.8%	11.7 ± 0.3%	7.3 ± 0.5%	4.8 ± 0.3%
SPiRL	18.5 ± 1.3%	15.1 ± 1.1%	9.7 ± 0.6%	12.8 ± 1.1%	8.9 ± 1.3%	6.1 ± 1.2%
SPiRL*	35.7 ± 4.1%	24.4 ± 3.9%	23.1 ± 3.4%	27.8 ± 3.5%	22.1 ± 4.1%	15.3 ± 2.7%
FIST	62.5 ± 7.0%	51.3 ± 6.9%	43.1 ± 9.3%	63.1 ± 4.6%	45.6 ± 4.4%	33.7 ± 5.9%
DCMRL	77.3 ± 6.2%	68.1 ± 8.9%	51.8 ± 4.5%	71.0 ± 4.8%	60.4 ± 7.2%	46.1 ± 8.0%
ICPAD (Ours)	<b>95.2 ± 4.9%</b>	<b>85.4 ± 7.1%</b>	<b>81.6 ± 7.4%</b>	<b>89.1 ± 10.3%</b>	<b>83.6 ± 7.5%</b>	<b>74.9 ± 8.8%</b>
<b>Evaluation on multi-stage Metaworld</b>						
DiffBC+FT	48.1 ± 2.7%	45.3 ± 2.5%	43.8 ± 1.6%	38.2 ± 2.2%	37.0 ± 0.9%	36.2 ± 2.1%
SPiRL	52.5 ± 0.8%	50.6 ± 2.3%	47.6 ± 1.0%	51.1 ± 0.9%	47.7 ± 1.0%	46.3 ± 2.2%
SPiRL*	58.2 ± 7.2%	56.1 ± 6.5%	54.7 ± 5.6%	57.0 ± 4.9%	54.3 ± 5.1%	54.1 ± 6.1%
FIST	89.4 ± 1.3%	78.3 ± 2.3%	68.9 ± 3.0%	85.2 ± 2.9%	75.5 ± 1.8%	64.3 ± 2.0%
DCMRL	<b>96.9 ± 0.9%</b>	90.0 ± 1.7%	81.7 ± 2.7%	91.0 ± 4.2%	81.9 ± 2.6%	76.1 ± 4.0%
ICPAD (Ours)	<b>96.9 ± 1.5%</b>	<b>94.8 ± 1.3%</b>	<b>90.2 ± 1.7%</b>	<b>92.8 ± 1.5%</b>	<b>90.6 ± 1.9%</b>	<b>85.2 ± 1.8%</b>

Table 1: Adaptation performance in Metaworld and multi-stage Metaworld w.r.t. environment dynamics: The performance in success rate is represented in 95% confidence intervals with 5 seeds. For each domain (Action noise and Wind), we test 3 different domain disparity settings (Low, Medium, High) between the source, which includes offline data and online source domains, and the target domain.

task. Thus, during policy learning and adaptation, the below dynamic prompting function  $\Psi_D$  is used.

$$\Psi_D(\mathcal{T}, \mathbf{s}^h) = \sum_{\hat{\tau}_i \sim \Phi_R(\mathcal{T})} \left[ \frac{w_i^{-1}}{\sum_{i=1}^m w_i^{-1}} \Psi_E(\hat{\mathbf{s}}_i, \hat{\mathbf{a}}_i) \right] \quad (13)$$

**Policy learning and adaptation.** The skill-based policy  $\pi$  is trained with the skill adapter  $\Phi_A$  and the domain encoder  $\Psi_E$  in the source domain, and it is regularized by the cross-domain skill prior  $\Phi_R$ . Similar to (Pertsch, Lee, and Lim 2021), for  $\mathcal{B}_R$  sampled from replay buffer  $\mathcal{D}_R$ , the learning objective includes Q value maximization and  $\Phi_R$ -regularization.

$$\max_{\pi} \mathbb{E}_{\mathcal{B}_R} [\bar{r}(s_t, z_t) - \lambda D_{\text{KL}}(\pi(\cdot|s_t), \Phi_R(\cdot|s_t))] \quad (14)$$

Here,  $\bar{r}(s_t, z_t)$  is the sum of the rewards obtained by executing skill  $z_t$  at state  $s_t$  and  $\lambda$  is a hyperparameter. In the source domain, we use the replay buffer for data  $\mathcal{T}$ .

In policy adaptation, given few-shot target domain data  $\mathcal{T}_T$ , the policy  $\pi$  (trained in the source domain) is capable of producing an action sequence by

$$\mathbf{a} = \Phi_A(s_t, \Psi_D(\mathcal{T}_T, \mathbf{s}^h), z_t \sim \pi(\cdot|s_t)) \quad (15)$$

for the target domain and state history  $\mathbf{s}^h = (s_{t-H}, \dots, s_t)$ .

## Experiments

We evaluate our framework in environments with diverse tasks and domains, including robotics manipulation in Metaworld (Yu et al. 2020), and autonomous driving with diverse routes in CALRA (Dosovitskiy et al. 2017).

**Baselines.** **DiffBC+FT** (Pearce et al. 2023) is a diffusion-based imitation learning baseline that incorporates fine-tuning. **SPiRL** (Pertsch, Lee, and Lim 2021) is a skill-based RL baseline for long-horizon environments. We assess two variations of SPiRL for few-shot adaptation: the standard

SPiRL which fine-tunes a skill-based policy and SPiRL\* which additionally fine-tunes the skill decoder (adapter). **FIST** (Hakhamaneshi et al. 2022) is a skill-based few-shot imitation method, enabling rapid adaptation through a semi-parametric approach to skill determination. **DCMRL** (He et al. 2024) is a state-of-the-art skill-based meta-RL method that enhances adaptability by decoupling the Gaussian context and the skill decoder to adjust to environmental contexts.

## Cross-domain Adaptation Performance

**Cross-domain w.r.t. environment dynamics.** To evaluate the adaptation performance across domain changes with varied environmental dynamics, we use the Metaworld and multi-stage Metaworld with additional action noise and wind settings. The Metaworld includes the multiple directional reach task, and the multi-stage Metaworld provides long-horizon scenarios for training agents in complex sequential sub-tasks such as closing a drawer and pressing a button. For the multi-stage Metaworld, we assess performance by measuring the success rate for each sub-task, calculated as the average ratio of successfully completed sub-tasks. For in-context policy adaptation, we use 5-shot demonstrations for each task. The domain disparity indicates the difference between the source, which includes offline data and online source domains, and target domains (i.e., differences in action noise or wind). Details of the domain disparity settings are in Appendix.

In Table 1, ICPAD consistently outperforms the baselines, demonstrating its superiority with an average success rate that is 14.0% higher than the closest competitor, DCMRL. SPiRL shows low performance due to its skill decoder’s limitations with offline data; SPiRL\* improves by fine-tuning more components but lacks stability and efficiency. FIST and DCMRL, while adapting skill-based policies, suffer significant performance drops of 20.1% and 22.6% respectively due to insufficient skill generalization as domain disparity

Domain factor	Vehicle embodiment			Vehicle embodiment + Weather		
	Low	Medium	High	Low	Medium	High
DiffBC+FT	21.4 ± 0.9%	17.6 ± 0.7%	15.1 ± 0.7%	19.7 ± 0.8%	16.9 ± 0.9%	13.2 ± 1.3%
SPiRL	26.5 ± 1.2%	22.0 ± 1.2%	20.2 ± 1.4%	24.3 ± 1.0%	20.7 ± 1.3%	17.2 ± 1.5%
SPiRL*	33.1 ± 1.8%	26.4 ± 0.8%	23.8 ± 2.1%	32.1 ± 1.5%	25.6 ± 1.4%	19.9 ± 1.7%
FIST	60.8 ± 2.6%	51.0 ± 5.7%	47.7 ± 1.6%	58.9 ± 2.6%	49.5 ± 1.9%	43.7 ± 2.7%
DCMRL	68.4 ± 2.2%	59.7 ± 1.3%	56.6 ± 2.7%	67.5 ± 3.2%	58.5 ± 3.3%	51.9 ± 3.2%
ICPAD (Ours)	<b>81.3 ± 2.8%</b>	<b>77.3 ± 4.2%</b>	<b>73.4 ± 2.0%</b>	<b>79.1 ± 1.8%</b>	<b>76.3 ± 2.7%</b>	<b>73.5 ± 2.0%</b>

Table 2: Adaptation performance in CARLA w.r.t. embodiment: The performance is normalized based on the maximum return by an expert policy in the target domain.

Domain factor	Horizon length		
	Low	Medium	High
DiffBC+FT	87.7 ± 2.8%	85.1 ± 3.5%	77.6 ± 4.7%
SPiRL	92.1 ± 1.2%	92.5 ± 1.1%	82.4 ± 1.6%
SPiRL*	97.8 ± 1.7%	94.5 ± 2.1%	86.0 ± 3.3%
FIST	<b>99.4 ± 0.5%</b>	95.2 ± 1.1%	94.2 ± 1.0%
DCMRL	99.1 ± 0.4%	96.3 ± 0.5%	94.9 ± 0.8%
ICPAD (Ours)	<b>99.4 ± 0.2%</b>	<b>97.8 ± 0.6%</b>	<b>96.6 ± 0.7%</b>

Table 3: Adaptation performance in multi-stage Metaworld w.r.t. horizon length.

increases. In contrast, our ICPAD shows only a slight performance degradation of 10.5%. ICPAD effectively employs cross-domain consistent learning for domain-agnostic prototype skills, which generalize skills across domains, and the domain-grounded skill adapter transforms these prototype skills into action sequences tailored to the specific characteristics of each domain. Using this approach, ICPAD consistently achieves better performance, even in diverse and unseen environmental dynamics.

**Cross-domain w.r.t. embodiment.** The CARLA simulator provides real-world scenarios for training agents in autonomous driving across diverse domains, including various vehicle embodiment and weather conditions. For multi-task scenarios, we configure different starting points and destinations to have various routes in navigation tasks. Table 2 presents the normalized returns in CARLA. ICPAD consistently surpasses the most competitive baseline, DCMRL, by 11.6% to 21.6%. Unlike FIST and DCMRL, which require direct adaptation of skill-based policies, the skill prompting approach allows ICPAD to simplify the complexity of the environment by enabling adaptation at the skill abstraction level. This facilitates the effective use of skill-based policies across various domains, even in complex environments.

**Cross-domain w.r.t. horizon length.** Table 3 shows the adaptation performance across domain changes on task horizon, presenting the success rates in the multi-stage Metaworld. All the skill-based baselines effectively address the changes in horizon length. Among them, ICPAD shows slightly better performance than the most competitive baseline DCMRL, by a margin of 0.3% to 1.7%. This robust performance can be attributed to dynamic domain prompting that selects appropriate prompts for each state, enabling effective adaptation to domains with extended horizons.

## Analysis

**Limited data availability for tasks.** Table 4 shows the average success rate of ICPAD with respect to data availability for the tasks from the target domain in the multi-stage Metaworld. In this test, we use only a selected percentage of tasks, each supported by 5-shot demonstrations. For tasks without data, we use data from the most similar task as a substitute. With only 8% of the tasks represented by demonstrations in the target domain, our ICPAD exhibits a slight performance decline of 6.8%, underscoring its robustness. In contrast, a 32.0% decrease is observed with DCMRL under the same conditions. This demonstrates the efficiency of our unified domain-wise policy adaptation, which utilizes skill prompting even with strictly limited data from the target domain.

Data availability	8%	20%	100%
FIST	22.0 ± 0.9%	47.2 ± 2.2%	75.5 ± 1.8%
DCMRL	49.9 ± 2.8%	70.3 ± 1.7%	81.9 ± 2.6%
ICPAD	<b>83.8 ± 2.1%</b>	<b>87.9 ± 1.7%</b>	<b>90.6 ± 1.9%</b>

Table 4: Impact of the data availability for tasks.

**Expansion of language-based prompts.** In real-world scenarios, there is a growing demand for an interactive human-agent system, where agents follow language instructions. To evaluate our framework in such a context, we utilize a dataset labeled with language instructions specific to each domain (e.g., The wind blows from the left). We integrate these language instructions as prompts in ICPAD, using an additional language encoder that is optimized by the domain encoder. For comparison, we implement L-DCMRL, which trains the language encoder with the context encoder. In Table 5, ICPAD maintains a superior performance compared to L-DCMRL. This specifies the ability of ICPAD to effectively utilize prompts from different modalities, achieving a similar performance to that obtained with sub-trajectories.

**Qualitative analysis for prototype skill embedding.** Figure 4 visualizes the prototype skill embedding generated by ICPAD for multi-stage Metaworld tasks. For different domains within the same task, the skill embeddings are closely paired together. This indicates that ICPAD effectively constructs shareable prototype skills, demonstrating that simply changing the prompt for the skill adapter allows seamless adaptation.

Domain factor	Instructions related to wind		
	Low	Medium	High
L-DCMRL	80.8 ± 1.1%	75.4 ± 1.5%	67.5 ± 1.4%
ICPAD	<b>90.4 ± 2.1%</b>	<b>87.1 ± 1.6%</b>	<b>82.9 ± 1.4%</b>

Table 5: Performance with language-based prompts in multi-stage Metaworld.

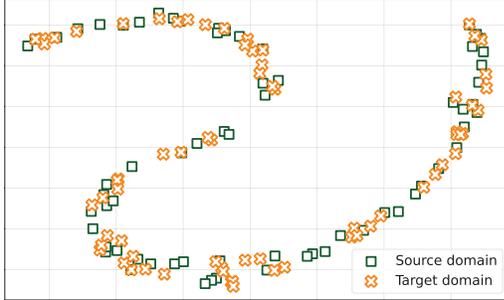


Figure 4: Qualitative analysis for prototype skill embedding.

## Ablation Study

We conduct ablation tests with the multi-stage Metaworld.

**Ablation for cross-domain skill diffusion.** In Table 6, to show the effect of the cross-domain skill diffusion, we test two variants of ICPAD; ICPAD-A does not use the cross-domain action consistency loss for the skill adapter  $\mathcal{L}^{\text{cross-A}}$  in (10), and ICPAD-D does not use diffusion models for the skill adapter. As shown, ICPAD outperforms these two variants, achieving a performance gain of 5.4% to 6.3%.

Domain factor	Wind		
	Low	Medium	High
ICPAD-A	79.6 ± 2.2%	76.4 ± 3.0%	53.9 ± 2.1%
ICPAD-D	87.4 ± 3.3%	84.1 ± 2.6%	78.9 ± 3.1%
ICPAD	<b>92.8 ± 1.5%</b>	<b>90.6 ± 1.9%</b>	<b>85.2 ± 1.8%</b>

Table 6: Ablation study for cross-domain skill diffusion.

**Ablation for dynamic domain prompting.** In Table 7, to show the effect of the dynamic domain prompting, we test two variants of ICPAD; ICPAD-Fix does not use a retrieval approach, but it uses the average of entire data for domain embedding. ICPAD-Avg does not use attention, but it uses the average of retrieved prompts for domain embedding. The results show that ICPAD outperforms these variants, demonstrating an average performance gain of 1.6% to 5.5%. Although the performance gap is minimal at lower levels of domain disparity, it becomes more pronounced as the disparity increases. This is because retrieval-based attention supports flexible adaptation to unseen domains.

## Related Works

**Policy adaptation.** In the application areas of autonomous systems where interaction with the target environment is limited or data from the target is scarce, policy adaptation methods have been investigated. Previous research primarily

Domain factor	Wind		
	Low	Medium	High
ICPAD-Fix	78.5 ± 2.7%	67.9 ± 1.8%	54.1 ± 2.9%
ICPAD-Avg	91.2 ± 1.7%	88.4 ± 1.4%	79.7 ± 2.6%
ICPAD	<b>92.8 ± 1.5%</b>	<b>90.6 ± 1.9%</b>	<b>85.2 ± 1.8%</b>

Table 7: Ablation study for dynamic domain prompting.

focused on searching a favorable initial weight (Finn, Abbeel, and Levine 2017; Nichol, Achiam, and Schulman 2018), learning policy-based task embeddings (Zintgraf et al. 2021), or adjusting rewards to mitigate domain differences (Eysenbach et al. 2021; Liu, Zhang, and Wang 2022). However, these approaches often encounter challenges in more complex, data-restricted environments, particularly when online interaction with the target is not allowed or when only limited expert data is available. In contrast, our focus is on addressing these data limitations in the target domain by integrating skill diffusion with cross-domain consistency learning, thus enabling in-context policy adaptation across domains.

**Consistency learning.** To achieve models with cross-domain capabilities, research is being conducted on extracting and transforming domain characteristics through the notion of consistency learning (Zhu et al. 2017; Chen et al. 2019; Wang et al. 2023). Particularly, in the RL literature, to facilitate policy transfer, CDIO (Raychaudhuri et al. 2021) and RL-CycleGAN (Rao et al. 2020) establish domain correspondence mappings via consistency learning. Our cross-domain skill diffusion incorporates consistency learning within the diffusion-based skill translation, aiming for a unified policy adaptation strategy based on prototype skills at a middle-tier level. Our work is the first to achieve in-context adaptation of skill-based policies, exploring the integration of skill diffusion and consistency learning.

**Skill-based RL.** To tackle complex tasks with challenging exploration, research efforts have been focused on incorporating short-horizon skills in hierarchical RL and multi-task RL (Eysenbach et al. 2018; Gregor, Rezende, and Wierstra 2016; Sharma et al. 2020). Recent works adopt offline skill imitation from task-agnostic expert datasets to accelerate policy learning (Pertsch, Lee, and Lim 2021; Nam et al. 2022; Pertsch et al. 2022). Our work also leverages skill imitation but distinguishes itself by concentrating on cross-domain skill learning and in-context policy adaptation.

## Conclusion

We presented the ICPAD framework to enable in-context policy adaptation across domains, effectively addressing the challenges posed by data limitations in the target domain. Central to the framework are the domain-agnostic prototype skills and their domain-specific translation via the diffusion-based skill adapter, achieved through consistency learning. The framework also employs dynamic domain prompting to enhance the performance of in-context policy adaptation. In our future work, we plan to extend the framework to accommodate multi-modal datasets, aiming to explore the semantic interpretability and alignment across significantly different tasks and domains for embodied control applications.

## Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT), (RS-2022-II220043 (2022-0-00043), Adaptive Personality for Intelligent Agents, RS-2022-II221045 (2022-0-01045), Self-directed multi-modal Intelligence for solving unknown, open domain problems, No. Rs-2020-II201821, ICT Creative Consilience Program, and RS-2019-II190421, Artificial Intelligence Graduate School Program (Sungkyunkwan University)), the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00213118), IITP-ITRC (Information Technology Research Center) grant funded by the Korea government (MIST) (IITP-2024-RS-2024-00437633, 10%), and by Samsung Electronics.

## References

- Chen, Y.-C.; Lin, Y.-Y.; Yang, M.-H.; and Huang, J.-B. 2019. Crdoco: Pixel-level domain transfer with cross-domain consistency. In *Proceedings of the 32nd IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1791–1800.
- Dosovitskiy, A.; Ros, G.; Codevilla, F.; Lopez, A.; and Koltun, V. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Conference on Robot Learning (CoRL)*, 1–16.
- Eysenbach, B.; Asawa, S.; Chaudhari, S.; Levine, S.; and Salakhutdinov, R. 2021. Off-dynamics reinforcement learning: Training for transfer with domain classifiers. In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*.
- Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2018. Diversity is all you need: Learning skills without a reward function. In *Proceedings of the 6th International Conference on Learning Representations (ICLR)*.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 1126–1135. PMLR.
- Gregor, K.; Rezende, D. J.; and Wierstra, D. 2016. Variational intrinsic control. *arXiv preprint arXiv:1611.07507*.
- Hakhamaneshi, K.; Zhao, R.; Zhan, A.; Abbeel, P.; and Laskin, M. 2022. Hierarchical few-shot imitation with skill transition models. In *Proceedings of the 10th International Conference on Learning Representations (ICLR)*.
- He, H.; Zhu, A.; Liang, S.; Chen, F.; and Shao, J. 2024. Decoupling Meta-Reinforcement Learning with Gaussian Task Contexts and Skills. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12358–12366.
- HO, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *Proceedings of the 34th Conference on Neural Information Processing System (NeurIPS)*.
- Hua, J.; Zeng, L.; Li, G.; and Ju, Z. 2021. Learning for a robot: Deep reinforcement learning, imitation learning, transfer learning. *Sensors*, 21(4): 1278.
- Liu, J.; Zhang, H.; and Wang, D. 2022. DARA: Dynamics-aware reward augmentation in offline reinforcement learning. In *Proceedings of the 10th International Conference on Learning Representations (ICLR)*.
- Nam, T.; Sun, S.-H.; Pertsch, K.; Hwang, S. J.; and Lim, J. J. 2022. Skill-based meta-reinforcement learning. In *Proceedings of the 10th International Conference on Learning Representations (ICLR)*.
- Nichol, A.; Achiam, J.; and Schulman, J. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*.
- Pearce, T.; Rashid, T.; Kanervisto, A.; Bignell, D.; Sun, M.; Georgescu, R.; Macua, S. V.; Tan, S. Z.; Momennejad, I.; Hofmann, K.; and Devlin, S. 2023. Imitating Human Behaviour with Diffusion Models. In *Proceedings of the 11th International Conference on Learning Representations (ICLR)*.
- Pertsch, K.; Desai, R.; Kumar, V.; Meier, F.; Lim, J. J.; Batra, D.; and Rai, A. 2022. Cross-domain transfer via semantic skill imitation. In *Proceedings of the 5th Conference on Robot Learning (CoRL)*.
- Pertsch, K.; Lee, Y.; and Lim, J. 2021. Accelerating reinforcement learning with learned skill priors. In *Proceedings of the 4th Conference on Robot Learning (CoRL)*, 188–204. PMLR.
- Rao, K.; Harris, C.; Irpan, A.; Levine, S.; Ibarz, J.; and Khansari, M. 2020. RL-CycleGAN: Reinforcement learning aware simulation-to-real. In *Proceedings of the 31st IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11157–11166.
- Raychaudhuri, D. S.; Paul, S.; Vanbaar, J.; and Roy-Chowdhury, A. K. 2021. Cross-domain imitation from observations. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 8902–8912. PMLR.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 815–823.
- Sharma, A.; Gu, S.; Levine, S.; Kumar, V.; and Hausman, K. 2020. Dynamics-aware unsupervised discovery of skills. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*.
- Shi, L. X.; Lim, J. J.; and Lee, Y. 2022. Skill-based model-based reinforcement learning. In *Proceedings of the 5th Conference on Robot Learning (CoRL)*.
- Wang, K.; Kim, D.; Feris, R.; and Betke, M. 2023. CDAC: Cross-domain attention consistency in transformer for domain adaptive semantic segmentation. In *Proceedings of the 17th IEEE/CVF International Conference on Computer Vision (ICCV)*, 11519–11529.
- Xu, M.; Shen, Y.; Zhang, S.; Lu, Y.; Zhao, D.; Tenenbaum, J.; and Gan, C. 2022. Prompting decision transformer for few-shot policy generalization. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, 24631–24645. PMLR.
- Yu, T.; Quillen, D.; He, Z.; Julian, R.; Hausman, K.; Finn, C.; and Levine, S. 2020. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning.

In *Proceedings of the 3rd Conference on Robot Learning (CoRL)*, 1094–1100. PMLR.

Zhao, W.; Queralta, J. P.; and Westerlund, T. 2020. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In *IEEE Symposium Series on Computational Intelligence (SSCI)*, 737–744. IEEE.

Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the 16th IEEE International Conference on Computer Vision (ICCV)*, 2223–2232.

Zhu, Z.; Lin, K.; Jain, A. K.; and Zhou, J. 2023. Transfer learning in deep reinforcement learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Zintgraf, L.; Schulze, S.; Lu, C.; Feng, L.; Igl, M.; Shiarlis, K.; Gal, Y.; Hofmann, K.; and Whiteson, S. 2021. Varibad: Variational bayes-adaptive deep rl via meta-learning. *Journal of Machine Learning Research*, 22(289): 1–39.