

Fast Adaptation with Behavioral Foundation Models

Anonymous authors
Paper under double-blind review

Keywords: Unsupervised Learning, Self Supervised learning, Zero shot RL, Adaptation, Finetuning

Summary

Unsupervised zero-shot reinforcement learning (RL) has emerged as a powerful paradigm for pretraining behavioral foundation models (BFMs), enabling agents to solve a wide range of downstream tasks specified via reward functions without additional test-time learning or planning. This is achieved by learning self-supervised task embeddings alongside corresponding near-optimal behaviors, and incorporating an inference procedure to directly retrieve the latent task embedding and associated policy for any given reward function. In this work, we demonstrate that existing unsupervised zero-shot RL pre-training methods discover a latent task embedding space containing more performant policies than those identified by their inference procedure, making them well-suited for fast adaptation. Motivated by this observation, we propose both actor-critic and actor-only fast adaptation strategies that search in the low-dimensional task-embedding space of the pre-trained BFM to rapidly improve the performance of its zero-shot policies on any downstream task. Notably, our approach mitigates the initial “unlearning” phase commonly observed when fine-tuning pre-trained RL models. We evaluate our fast adaptation strategies on top of four state-of-the-art zero-shot RL methods in multiple navigation and locomotion domains. Our results show that they achieve 10-40% improvement over their zero-shot performance in only a few episodes, outperforming existing baselines.

Contribution(s)

1. We empirically investigate the task-representation space learned by a family of unsupervised zero-shot RL methods and show that it contains policies achieving significantly higher returns than the one output by the zero-shot inference procedure.
Context: Prior works in zero-shot RL (Touati et al., 2023; Park et al., 2024; Agarwal et al., 2024) implicitly assume that zero-shot inference is the optimal way to prompt a pre-trained model for behaviors optimizing tasks specified by reward functions. We challenge such an assumption and show that this is not the case.
2. We propose two fast-adaptation algorithms: a) Residual Latent Adaptation (ReLA), an approach that optimizes for a policy in the BFM’s task-representation space by training an additional smaller critic to estimate the cumulative reward not captured by the pre-trained BFM. b) Lookahead Latent Adaptation (LoLA), a computationally efficient approach that leverages policy gradients with lookahead returns without updating the pre-trained critic.
Context: Prior approaches to adaptation either fine-tune the entire pre-trained critic and perform policy optimization in the action space (Nair et al., 2020; Nakamoto et al., 2023), or learn policy residuals (Silver et al., 2018; Johannink et al., 2019; Rana et al., 2023).
3. We evaluate our approaches on top of four state-of-the-art zero-shot RL methods in multiple navigation and locomotion domains, and show that they achieve 10-40% improvement over their zero-shot performance. Furthermore, we observe that our approach LoLA avoids the initial “unlearning” phase commonly observed in the literature.
Context: Prior approaches for fine-tuning RL models without retaining training data (Luo et al., 2023; Zhou et al., 2024) observe a sharp decrease in performance due to distribution shift.

Fast Adaptation with Behavioral Foundation Models

Anonymous authors

Paper under double-blind review

Abstract

Unsupervised zero-shot reinforcement learning (RL) has emerged as a powerful paradigm for pretraining behavioral foundation models (BFMs), enabling agents to solve a wide range of downstream tasks specified via reward functions in a zero-shot fashion, i.e., without additional test-time learning or planning. This is achieved by learning self-supervised task embeddings alongside corresponding near-optimal behaviors and incorporating an inference procedure to directly retrieve the latent task embedding and associated policy for any given reward function. Despite promising results, zero-shot policies are often suboptimal due to errors induced by the unsupervised training process, the embedding, and the inference procedure. In this paper, we focus on devising *fast adaptation* strategies to improve the zero-shot performance of BFMs in few steps of online interaction with the environment, while avoiding any performance drop during the adaptation process. Notably, we demonstrate that existing BFMs learn a set of skills containing more performant policies than those identified by their inference procedure, making them well-suited for fast adaptation. Motivated by this observation, we propose both actor-critic and actor-only fast adaptation strategies that search in the low-dimensional task-embedding space of the pre-trained BFM to rapidly improve the performance of its zero-shot policies on any downstream task. Notably, our approach mitigates the initial “unlearning” phase commonly observed when fine-tuning pre-trained RL models. We evaluate our fast adaptation strategies on top of four state-of-the-art zero-shot RL methods in multiple navigation and locomotion domains. Our results show that they achieve 10-40% improvement over their zero-shot performance in a few tens of episodes, outperforming existing baselines.

1 Introduction

Unsupervised (or self-supervised) pre-training has emerged as one of the key ingredients behind the recent breakthroughs in computer vision and language modeling (e.g., [Radford et al., 2019](#); [Devlin et al., 2019](#); [Touvron et al., 2023](#); [Caron et al., 2021](#)). This technique allows utilizing large datasets of unlabeled data samples to learn generalizable representations that can be later fine-tuned for various downstream applications ([Zhai et al., 2023](#); [Brown et al., 2020](#); [Driess et al., 2023](#)). For instance, language models are pre-trained on internet-scale data with a next-token prediction objective and later fine-tuned for desired applications using high-quality examples. How to transpose this approach to reinforcement learning (RL) to train agents that can efficiently solve sequential decision-making problems is an open research question of paramount importance. Going beyond the tabula-rasa paradigm of classic RL requires an unsupervised pre-training objective and the ability to efficiently fine-tune or adapt pre-trained representations for downstream tasks. Recent developments in unsupervised RL propose various objectives to learn a repertoire of skills on top of reward-free data from the environment ([Gregor et al., 2016](#); [Wu et al., 2018](#); [Hansen et al., 2019](#); [Liu & Abbeel, 2021](#); [Eysenbach et al., 2018](#); [Zahavy et al., 2022](#); [Park et al., 2023](#)). Some of these methods are named “zero-shot”, in the sense that they additionally provide a procedure to infer a performant policy for any given task specified by reward functions ([Touati et al., 2023](#); [Park et al., 2024](#); [Agarwal et al., 2024](#); [Cetin et al., 2024](#)), demonstrations ([Pirota et al., 2024](#); [Tirinzoni et al., 2025](#)), or

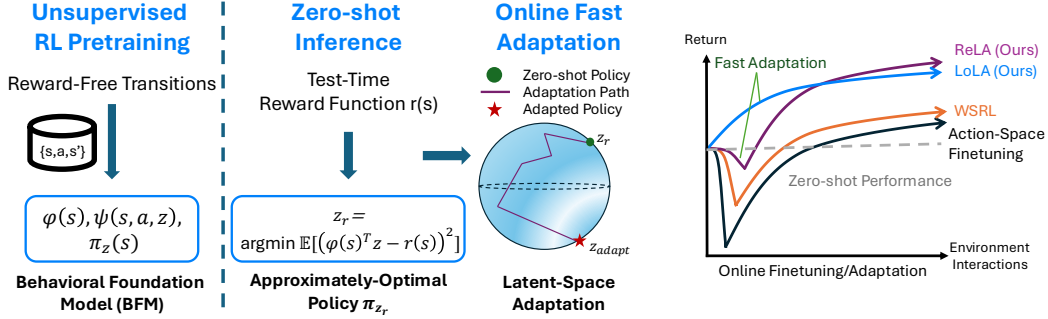


Figure 1: **Overview of our method:** Unsupervised zero-shot RL methods provide us with an initial policy π_{z_r} ; we propose a way to leverage the latent space of learned policies as well as the pre-trained critic to rapidly adapt and improve π_{z_r} on few task-specific environment interactions. Right: Illustrative summary of our results.

41 videos/language (Sikchi et al., 2024). The resulting pre-trained agents are commonly referred to as
 42 Behavioral Foundation Models (BFMs, Pirotta et al., 2024; Tirinzoni et al., 2025).

43 Zero-shot methods commonly pre-train two components: (1) a *state representation* $\varphi : \mathcal{S} \rightarrow \mathbb{R}^d$
 44 that embeds state observations $s \in \mathcal{S}$ into a d -dimensional vector $\varphi(s)$, and (2) a space $\{\pi_z\}$ of
 45 *policies* parameterized by a latent vector $z \in \mathbb{R}^d$. The representation φ defines the set of all linear
 46 reward functions in φ , i.e., $\tilde{r}_z(s) = \varphi(s)^T z$ for all $z \in \mathbb{R}^d$, which in turn is used as a *self-supervised*
 47 objective function for the policy space: for each $z \in \mathbb{R}^d$, the policy π_z is trained to be approximately
 48 optimal for the reward \tilde{r}_z . Given a reward function $r(s)$ at test time, a zero-shot policy π_{z_r} can be
 49 obtained by projecting r onto the pre-trained state features φ through linear regression on top of the
 50 training data, hence approximating $r(s) \simeq \varphi(s)^T z_r$.

51 Although this inference method has proven effective in producing reasonable policies, it suffers
 52 from two main limitations yielding sub-optimal performance. First, the embedding φ is learned
 53 using unsupervised losses encoding inductive biases¹ that may not be suitable for the downstream
 54 tasks of interest. As a result, the projection of the reward function onto φ may remove crucial aspects
 55 of the task specification thus preventing from finding the optimal policy for the original reward. In an
 56 extreme scenario, if a reward function lies in the orthogonal subspace of the features’ linear span, its
 57 projection onto these features becomes zero, making it uninformative. Second, BFMs are typically
 58 trained on task-agnostic datasets that may have poor coverage of the rewarding states relevant to the
 59 specific task. This limitation can result in zero-shot inference failing to accurately represent these
 60 states and ultimately hinder the learning of a good policy.

61 While the suboptimality of unsupervised pre-training of large models is somewhat unavoidable, it
 62 is natural to wonder whether these limitations can be overcome once a downstream reward function
 63 is given and the agent has online access to the environment. In this paper we focus on devising *fast*
 64 *adaptation* strategies that improve zero-shot performance of BFMs **1) rapidly**, i.e., in a handful of
 65 online episodes, and **2) monotonically**, i.e., avoiding any performance drop during the adaptation
 66 process. This motivates the main question of this work:

67 *Does the policy space of a pre-trained BFM contain better behaviors than those returned by*
 68 *zero-shot inference? If so, can we retrieve them with few task-specific environment interactions?*

69 To address this question, we propose searching over the latent space \mathcal{Z} using a limited number of
 70 online task-specific interactions with the environment (cf. Figure 1). We introduce two algorithms
 71 that leverage the latent space and pre-trained components from BFMs to enable *fast adaptation* of
 72 their zero-shot policies: (1) Residual Latent Adaptation (ReLA), an off-policy actor-critic approach
 73 that trains a small *residual critic* to compensate for the reward projection errors, and Lookahead

¹For instance, some methods rely on low-rank assumptions in the policy dynamics (Touati & Ollivier, 2021; Agarwal et al., 2024), while others focus only on goal-reaching behaviors (Park et al., 2024)




Figure 2: Performance comparison of zero-shot policy vs adapted policy in the BFM’s latent space after 200 episodes. The shaded region shows the improvement of the adapted policies averaged across tasks.

74 Latent Adaptation (LoLA), a hybrid actor-only approach that combines on-policy optimization while
 75 bootstrapping the frozen critic from the pre-trained BFMs.

76 We perform an extensive empirical evaluation on 5 domains with a total of 64 tasks spanning low-
 77 dimensional and high-dimensional problems with increasing complexity, including a whole-body
 78 humanoid control problem with a wide range of 45 diverse reward-based behaviors. We demon-
 79 strate the effectiveness of our proposed algorithms on four state-of-the-art BFMs: FB (Touati &
 80 Ollivier, 2021), HILP (Park et al., 2024), PSM (Agarwal et al., 2024) and FB-CPR (Tirinzoni et al.,
 81 2025). In particular, we answer the above question affirmatively: our fast adaptation algorithms
 82 achieve 10-40% improvement over the BFMs zero-shot performance in only a few episodes (Fig-
 83 ure 2), while outperforming existing baselines. Moreover, we show that LoLA avoids any initial
 84 drop of performance, a phenomenon commonly observed by numerous prior works on fine-tuning
 85 RL policies (Nair et al., 2020; Nakamoto et al., 2023; Luo et al., 2023; Zhou et al., 2024).

86 2 Preliminaries

87 **Markov decision process.** We consider a reward-free Markov decision process (MDP) (Puter-
 88 man, 2014; Sutton & Barto, 2018) which is defined as a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, d_0, \gamma)$, where \mathcal{S} and
 89 \mathcal{A} respectively denote the state and action spaces, P denotes the transition kernel with $P(s'|s, a)$
 90 indicating the probability of transitioning from s to s' by taking action a , d_0 denotes the initial state
 91 distribution and $\gamma \in (0, 1)$ specifies the discount factor. A policy π is a function $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$
 92 mapping a state s to probabilities of action in \mathcal{A} . We denote by $\Pr(\cdot | s, a, \pi)$ and $\mathbb{E}[\cdot | s, a, \pi]$ the
 93 probability and expectation operators under state-action sequences $(s_t, a_t)_{t \geq 0}$ starting at (s, a) and
 94 following policy π with $s_t \sim P(\cdot | s_{t-1}, a_{t-1})$ and $a_t \sim \pi(\cdot | s_t)$. Given any reward function
 95 $r : \mathcal{S} \rightarrow \mathbb{R}$, the Q-function of π for r is $Q^\pi_\pi(s, a) := \sum_{t \geq 0} \gamma^t \mathbb{E}[r(s_{t+1}) | s, a, \pi]$.

96 **Successor measures and features.** The *successor measure* (Dayan, 1993; Bluer et al., 2021) of
 97 state-action (s, a) under a policy π is the (discounted) distribution of future states obtained by taking
 98 action a in state s and following policy π thereafter:

$$M^\pi(X | s, a) := \sum_{t \geq 0} \Pr(s_{t+1} \in X | s, a, \pi) \quad \forall X \subset \mathcal{S}. \quad (1)$$

99 Importantly, successor measures disentangle the dynamics of the MDP and the reward function: for
100 any reward r and policy π , the Q-function can be expressed linearly as $Q_r^\pi = M^\pi r$.

101 Given a feature map $\varphi : \mathcal{S} \rightarrow \mathbb{R}^d$ that embeds states into a d -dimensional space, the *successor*
102 *features* (Barreto et al., 2017) is the expected discounted sum of features:

$$\psi^\pi(s, a) := \sum_{t \geq 0} \gamma^t \mathbb{E}[\varphi(s_{t+1}) \mid s, a, \pi]. \quad (2)$$

103 Successor features and measures are related: by definition, $\psi^\pi(s, a) = \int_{\mathcal{S}} M^\pi(ds' \mid s, a) \varphi(s')$. For
104 any reward function in the linear span of φ , i.e., $r(s) = \omega^\top \varphi(s)$ where ω is a weight vector in \mathbb{R}^d ,
105 the Q-function can be expressed compactly as $Q_r^\pi(s, a) = \omega^\top \psi^\pi(s, a)$.

106 **Behavioral foundation models.** A *behavioral foundation model*, for a given MDP, is an agent
107 that can be trained in unsupervised fashion using reward-free transitions and yet can produce ap-
108 proximately optimal policies for a large class of reward functions r specified at test time, without
109 performing additional learning or planning. In this work, we focus on zero-shot RL agents that are
110 based on successor features and forward-backward representations.

111 *Universal successor features* (USFs) (Borsa et al., 2018) provide a generic framework for zero-shot
112 RL. Given a feature map φ , USFs learn the successor features of a particular family of policies π_z
113 parameterized by latent variables $z \in \mathcal{Z} \subset \mathbb{R}^d$:

$$\psi(s, a, z) = \mathbb{E}[\sum_{t \geq 0} \gamma^t \varphi(s_{t+1}) \mid s, a, \pi_z], \quad \pi_z(s) = \arg \max_a \psi(s, a, z)^\top z. \quad (3)$$

114 At test time, once a reward function r is specified, a reward-maximizing policy is inferred
115 by performing a linear regression of r onto the features φ . In particular, we estimate $z_r =$
116 $\arg \min_z \mathbb{E}_{s \sim \rho}[(r(s) - \varphi(s)^\top z)^2] = \mathbb{E}_{s \sim \rho}[\varphi(s) \varphi(s)^\top]^{-1} \mathbb{E}_{s \sim \rho}[\varphi(s) r(s)]$ where ρ is some dataset
117 distribution over states. Then we return the pre-trained policy π_{z_r} . This policy is guaranteed to be
118 optimal if the reward is in the linear span of the features φ (Borsa et al., 2018). Although USF is
119 a generic framework, it requires specifying a training criterion to learn the basic features φ . Touati
120 et al. (2023) compare several choices of unsupervised representation learning objectives across vari-
121 ous empirical problems. In this work, we focus on two recent state-of-the-art feature learning meth-
122 ods for zero-shot RL: *Hilbert representations* (HILP) (Park et al., 2024) and *proto successor mea-*
123 *sures* (PSM) (Agarwal et al., 2024). HILP constructs features φ such that the distance $\|\varphi(s) - \varphi(s')\|$
124 between a state pair (s, s') encodes the optimal value function of reaching the state s' starting at s .
125 PSM proposes to build the features φ by learning an *affine decomposition* of the successor measure
126 for a discrete codebook of policies, i.e., $M^{\pi_u}(ds' \mid s, a) / \rho(ds') \approx \phi(s, a)^\top \varphi(s') w(u) + b(s, a, s')$,
127 where ϕ, w and b are vector-valued functions and where π_u is a deterministic policy that outputs an
128 action in state s as a realization of the uniform distribution, determined by the random seed u .

129 *Forward-backward representations* (FB) (Touati & Ollivier, 2021) provide an alternative framework
130 for zero-shot RL. Unlike USFs which use two separate criteria to learn features and their successor
131 features, FB avoid the state featurization step and employ a single objective to learn a finite-rank de-
132 composition of the successor measure for various policies. Namely, FB pre-train two representations
133 $F : \mathcal{S} \times \mathcal{A} \times \mathcal{Z} \rightarrow \mathbb{R}^d$ and $B : \mathcal{S} \rightarrow \mathbb{R}^d$ such that:

$$F(s, a, z)^\top B(s') \rho(ds') \approx M^{\pi_z}(ds' \mid s, a), \quad \pi_z(s) = \arg \max_a F(s, a, z)^\top z. \quad (4)$$

134 FB representations are related to USFs, as $F(s, a, z)$ represents the successor features of
135 $\mathbb{E}_{s \sim \rho}[B(s) B(s)^\top]^{-1} B(s)$ (Touati et al., 2023). In the sequel, to standardize the notations with
136 the USFs, we will denote $\psi(s, a, z) = F(s, a, z)$ and $\varphi(s) = \mathbb{E}_{s \sim \rho}[B(s) B(s)^\top]^{-1} B(s)$.

137 *Forward-Backward representations with Conditional Policy Regularization* (FB-CPR) (Tirinzoni
138 et al., 2025) is an online variant of FB that grounds the unsupervised policy learning toward im-
139 itating observation-only unlabeled behaviors.

3 Fast Adaptation for Behavioral Foundation Models

In this section, we introduce our two approaches for fast adaptation of pre-trained BFM: an off-policy actor-critic algorithm (Section 3.1), and a hybrid on-policy actor-only algorithm (Section 3.2).

3.1 ReLA: Residual Latent Adaptation

Given a reward function r , ReLA begins with the latent variable $z = z_r$ inferred by the zero-shot procedure and uses an off-policy actor-critic approach to gradually update z towards better performance. The overall algorithm uses a standard online training procedure, interleaving between critic and actor updates (as described below), while gathering reward-labeled transitions in a replay buffer $\mathcal{D}_{\text{online}}$ through online interactions with the environment.

Residual critic learning. Instead of training a critic from scratch to model the Q-function of the policy π_z currently being learned for the reward r , ReLA uses a residual critic to correct for the reward projection error. This is made possible by the following decomposition:

$$\begin{aligned} Q_r^{\pi_z}(s, a) &= Q_{\varphi^\top z_r}^{\pi_z}(s, a) + Q_{r - \varphi^\top z_r}^{\pi_z}(s, a) \\ &= \psi(s, a, z)^\top z_r + Q_{r - \varphi^\top z_r}^{\pi_z}(s, a) \end{aligned} \quad (5)$$

where the last equality holds because ψ is pre-trained to estimate the successor features of φ and the projected reward $\varphi^\top z_r$ lies in the span of φ . Consequently, ReLA considers a network $Q^{\text{residual}}(s, a; \theta)$ parametrized by weights θ and trains it via off-policy TD learning so that $\psi(s, a, z_r)^\top z_r + Q^{\text{residual}}(s, a; \theta)$ approximates the Q-function $Q_r^{\pi_z}(s, a)$, while keeping the *base Q-function* $\psi(s, a, z_r)^\top z_r$ frozen. In practice, we shall use much smaller networks for the residual critic than for the pre-trained successor features, with the main intuition being that we only need to compensate for some projection error. For a more in-depth treatment of the Q-function decomposition we refer the readers to Appendix 6.1.

Latent actor update. ReLA updates the latent variable z using standard policy-gradient ascent, with the key difference being that the gradient is computed only with respect to z , while keeping the pre-trained actor parameters fixed,

$$\nabla_z \mathbb{E}_{s \sim \mathcal{D}_{\text{online}}} [\psi(s, \pi_z(s), z_r)^\top z_r + Q^{\text{residual}}(s, \pi_z(s); \theta)], \quad (6)$$

The main advantage over optimizing the whole actor network is that we only need to search in a low-dimensional space (in practice, z has in the order of hundreds of components, while the actor network of a BFM has in the order of millions of parameters).

3.2 LoLA: Lookahead Latent Adaptation

Although ReLA can take advantage of off-policy data collected in the replay buffer, it requires learning an additional residual network. Therefore, ReLA demands a certain budget of transitions and updates to mitigate the distribution shift issue (Luo et al., 2023) when learning the Q-function, which may impede improvements during the very early stages of adaptation. On the other hand, a purely on-policy approach will require rolling out entire trajectories under the current policy to estimate Monte Carlo returns $\sum_{t=0}^T \gamma^t r_t$ (where T is the episode length), and thus incur many environment interactions in the process. Alternatively, we propose **Lookahead Latent Adaptation** algorithm (LoLA) that uses fixed-horizon on-policy rollouts with a frozen terminal value function obtained from the BFM. LoLA parameterizes a policy over the latent space as a normal distribution $\pi_{\mu, \sigma} = \mathcal{N}(\mu, \sigma)$ with trainable mean μ (initialized with $\mu = z_r$), and fixed diagonal covariance σ . The pre-trained successor features from BFM are used compute the estimate of a terminal value-function, thus estimating the n -step lookahead return of policy π_z starting from state s_0 as $R^n(s_0, z) = \sum_{t=0}^{n-1} \gamma^t r(s_{t+1}) + \gamma^n \psi(s_{n+1}, \pi_z(s_{n+1}), z)^\top z_r$.

Moreover, to further improve learning, LoLA incorporates the variance reduction strategy of leave-one-out baseline (Kool et al., 2019). This baseline has recently been shown to be empirically effective for fine-tuning large language models (Ahmadian et al., 2024). This leads to the following final gradient estimate²:

$$\mathbb{E}_{s_0 \sim \nu} \left[\frac{1}{k} \sum_{i=1}^k \left(R(s_0, z_i) - \frac{1}{k-1} \sum_{\substack{j=1 \\ j \neq i}}^k R(s_0, z_j) \right) \nabla_{\mu} \log \pi_{\mu, \sigma}(z_i) \right] \text{ for } z_1, \dots, z_k \sim \pi_{\mu, \sigma}(\cdot) \quad (7)$$

where s_0 is sampled from the distribution ν defined as mixture between the environment’s initial distribution d_0 and the online replay buffer distribution $\mathcal{D}_{\text{online}}$. For each sampled starting state s_0 , we sample k latent variables $\{z_i\}_{i \in [k]} \sim \pi_{\mu, \sigma}$ and generate k trajectories of length n $(s_0^{(i)}, a_0^{(i)}, s_1^{(i)}, a_1^{(i)}, \dots, s_n^{(i)})$ by following the policy π_{z_i} . Computing the gradient requires the ability to reset of any state in support of distribution ν , which includes the states encountered during online adaptation.

4 Experimental Results

The goal of our experiments is to study how well latent policy adaptation works on top of existing BFM. We perform several ablations to understand the efficacy of the proposed methods and evaluate our design choices. Precisely, 1) Can we find better policies by online latent policy adaptation compared to the zero-shot policies? Or, equivalently, is the latent policy space easy to search over? 2) How important is to leverage BFM’s properties (e.g., Q-function estimate, zero-shot policy initialization)? 3) What are the critical limitations of the zero-shot inference process?

Experimental setup. We investigate these questions by leveraging 4 different BFM: FB, HILP, PSM and FB-CPR. While FB, HILP and PSM are trained offline, FB-CPR learns through online environmental interactions and it is regularized towards expert trajectories. We consider four environments from the DeepMind Control suite (Tassa et al., 2018) and train the BFM on an exploratory dataset obtained from ExoRL (Yarats et al., 2022).³ Further, we leverage the FB-CPR model released by Tirinzoni et al. (2025) for the HumEnv environment, a high-dimensional humanoid agent. Overall, we consider 7 tasks for Pointmass, 4 for Cheetah, 4 for Quadruped, 4 for Walker, and 45 tasks for HumEnv. Detailed information about the pre-training phase can be found in Appendix 8.

Protocol and baselines. While the paper focuses on adaptation in the latent policy space, we also investigate the common class of approaches for fine-tuning in action space (i.e., updating all policy network parameters) with zero-shot initialization (Nair et al., 2020; Nakamoto et al., 2023). In particular, we consider a TD3-based algorithm (Fujimoto et al., 2018) that train a critic from scratch and an actor initialized using the zero-shot policy (TD3 (I)).⁴ Since collecting a few on-policy trajectories before starting updating the critic and the actor proved to be effective for offline to online adaptation, a strategy called Warm-Start RL (WSRL) (Zhou et al., 2024), we additionally consider this component for action-based algorithms. We further ablate several design choices (e.g., zero-shot initialization, bootstrapped critic) in Section 4.2 and, in Appendix 9 we report variations of our algorithms that operate by directly updating the parameters of the policy. See Table 1 for a complete list of algorithm variations.

We use a comparable architecture and hyperparameter search for all algorithms. For each BFM, we report the performance on the set of hyperparameters that performed best across all tasks and domains. We train all the online adaptation algorithms for 300 episodes and we use 5 seeds for each experiment. Evaluation is done by averaging results over 50 episodes. We also use TD3 as base off-policy algorithm for implementing ReLA. When using residual critic we use a small 2-layers MLP with hidden dimension 64, while when we learn the critic from scratch we use a 2-layers MLP with hidden dimension 1024. The policy has always the same size as the BFM policy. We provide further implementation details in Appendix 8.

²In practice we work with z normalized on hypersphere using projected gradient descent

³We consider the dataset collected by running RND (Burda et al., 2019).

⁴We also tested vanilla RLOO (Kool et al., 2019) but did not get good results and decided not to report it.

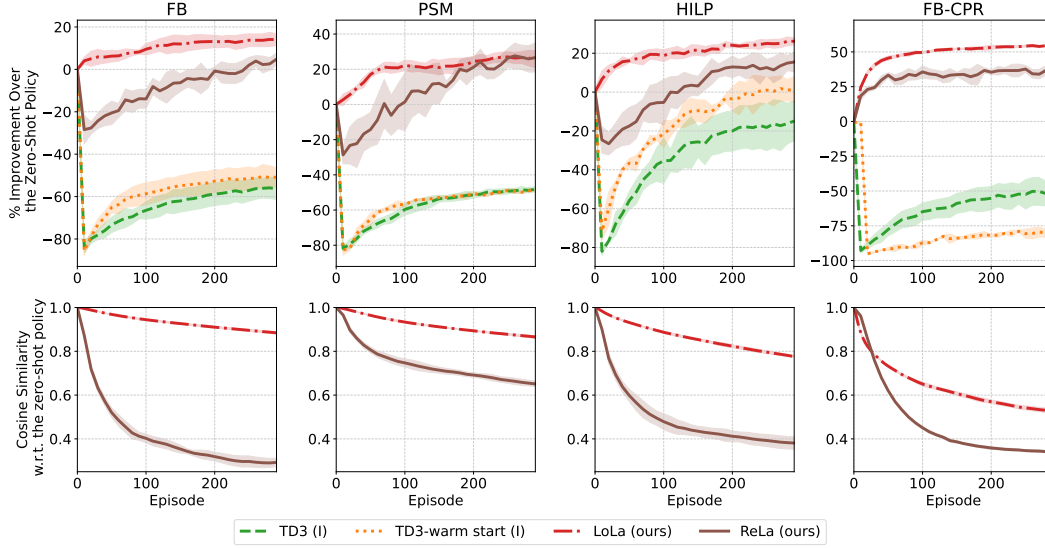


Figure 3: **Top:** Performance improvement w.r.t. the zero-shot policy for different online fast adaptation methods and BFM domains. TD3(I) denotes standard action-based TD3 with zero-shot policy initialization, our methods are as described in Section 3. **Bottom:** Cosine similarity between the zero-shot policy z_r and the learned policy z for the methods working in the latent policy space. We report mean and standard deviation over 5 seeds. Results are averaged over 19 tasks for FB, PSM, HILP and 45 tasks for FB-CPR.

224 4.1 Do ReLA and LoLA enable fast adaptation?

225 Figure 3 (top) shows our aggregated results across tasks for each domain: Latent policy adaptation
 226 leads to performance improvements w.r.t. the zero-shot policy in the range of 10-30% for DMC
 227 environment and 40-50% for HumEnv. Compared to Figure 2, these results show that significant
 228 improvements are already obtained in few online episodes. For example, LoLA leads to about 10%
 229 (resp. 40%) improvement for DMC (resp. for HumEnv) in only 20 episodes. These results show that
 230 i) the space of policies learned by the BFMs contains better policies than the one inferred by the zero-
 231 shot procedure and ii) such a space can be easily navigated using gradient-based approaches. While
 232 both ReLA and LoLA provide significant performance improvements, LoLA is the only method to
 233 achieve monotonic performance improvement across the board. As we can see from the per-task
 234 visualization in Appendix 9, the non-monotonic performance of ReLA is mostly due to the fact
 235 that the method incurs a noticeable catastrophic forgetting in the `pointmass` environment where
 236 TD3-based methods seem to struggle in the online setting, probably due to exploration issues. As
 237 a result of training purely on online samples, critic learning in ReLA undergoes distribution shift
 238 which has been investigated to lead to initial unlearning (Zhou et al., 2024) whereas LoLA skips the
 239 critic learning step entirely. In addition, LoLA exploits a privileged information compared to ReLA,
 240 the ability to reset the environment to any arbitrary state, which further contributes in stabilizing and
 241 speeding up the learning process (see e.g., Mhammedi et al., 2024).

242 **How does the adapted policy evolve in latent space?** To try to better understand the learning
 243 dynamics of ReLA and LoLA we report the cosine similarity between the adapted z and the zero-shot
 244 policy z_r in Figure 3 (bottom). ReLA deviates much more in the latent space from the initial zero-
 245 shot policy than LoLA. This fast and significant change is associated with the drop in performance.
 246 On the other hand, despite the high learning rate (we found 0.1 or 0.05 to be the best based on the
 247 BFM), LoLA remains closer to the zero-shot policy. A potential cause for the significant change in
 248 ReLA may be difficulties in critic learning associated with distribution shift, which can impact policy
 249 directly. This visualization also shows that while converging to different policies, the performance

of ReLA and LoLA is comparable after 300 episodes in the DMC environments. This reveals that policies with similar performance may be associated to with different latent vectors z .

When looking at the baselines, we can notice that all action-space adaptation algorithms suffer a much more significant drop compared to latent policy adaptation. The performance gap between action-based and latent policy adaptation becomes even larger when looking at FB-CPR. In this case, all action-based algorithms completely unlearn in a few steps and are not able to rapidly recover. We think this is due to the large dimensionality of observation space, action space and policy model that lead to a much more complicated optimization problem.⁵ On the other hand, in contrast to other BFMs, FB-CPR does not suffer any initial performance drop when using ReLA. Indeed, all latent policy adaptation algorithms (see Appendix 9 for additional experiments) achieve monotonic performance improvement, stressing even more that structured search in the latent policy space may be simpler than finetuning the whole policy in high-dimensional problems. This may be due to the fact that FB-CPR is the only BFM that is pre-trained with online environmental interactions, a setting that may reduce the distribution shift between pretraining and adaptation. Finally, the performance improvement due to the latent policy adaption is much more significant in this domain. The reason may reside in the critic training objective of FB-CPR; indeed FB-CPR uses a discriminator-based loss to regularize the policy space towards expert demonstrations. This may prevent the zero-shot inference to correctly identify the best policy for the task, while online adaptation seems to better search the policy space.

Finally, we would like to report an observation about the computational efficiency. On our hardware, LoLA runs at $\approx 157x$ the FPS of ReLA and other adaptation approaches. Specifically, ReLA runs at ≈ 14 FPS, and LoLA runs at $\approx 2,200$. This gaps presumably comes from the fact that ReLA needs to backpropagate gradient through the BFM estimated value function and policy both in the critic and actor updates, while LoLA has just a single actor update. The computational efficiency of LoLA along with its observed near-monotonic improvement for adaptation makes it appealing in practice.

4.2 What components are critical for fast adaptation?

In this section we assess the importance of leveraging BFM properties for fast online adaptation. We focus on ablating the need of *i*) zero-shot initialization and *ii*) BFM value function estimate, i.e., using a residual critic for ReLA and the bootstrapped Q-function for LoLA. We focus on the very early steps of the training to better inspect the results. Ablation variants are concisely shown in Table 1 for reference, and results are reported in Figure 4.

When *zero-shot initialization* is disabled not only the performance starts lower but also take significantly longer to match the baseline’s returns (if they match at all). Unsurprisingly, zero-shot initialization helps in the search process. Leveraging the BFM’s value function estimate does not hurt and often helps in reducing the initial performance drop. Looking at LoLA, BFM bootstrapping helps only marginally in all the domains. We believe that this is due to the small discount factor and large lookahead

Algorithm	Zero-Shot Policy Init.	Residual Critic([†])/ Bootstrapped Return(⁺)	Critic Trained from scratch	Search space
LoLA	✓	✓(⁺)		z
LoLA (no-I)		✓(⁺)		z
LoLA (no-R)	✓		✓	z
LoLA (no-I, no-R)			✓	z
ReLA	✓	✓([†])		z
ReLA (no-I)		✓([†])		z
ReLA (no-R)	✓		✓	z
ReLA (no-I, no-R)			✓	z
TD3-z			✓	z
TD3 (I)	✓		✓	a
TD3-warm-start(I)	✓		✓	a
TD3-warm-start(I, R)	✓	✓([†])		a

Table 1: Summary of algorithm variations. Here, search space z indicates latent policy adaptation via the policy space $\{\pi_z\}$ constructed by the BFM, while a denotes fine-tuning in action space.

⁵A way to address this problem may be through policy regularization but this is outside the scope of this paper.

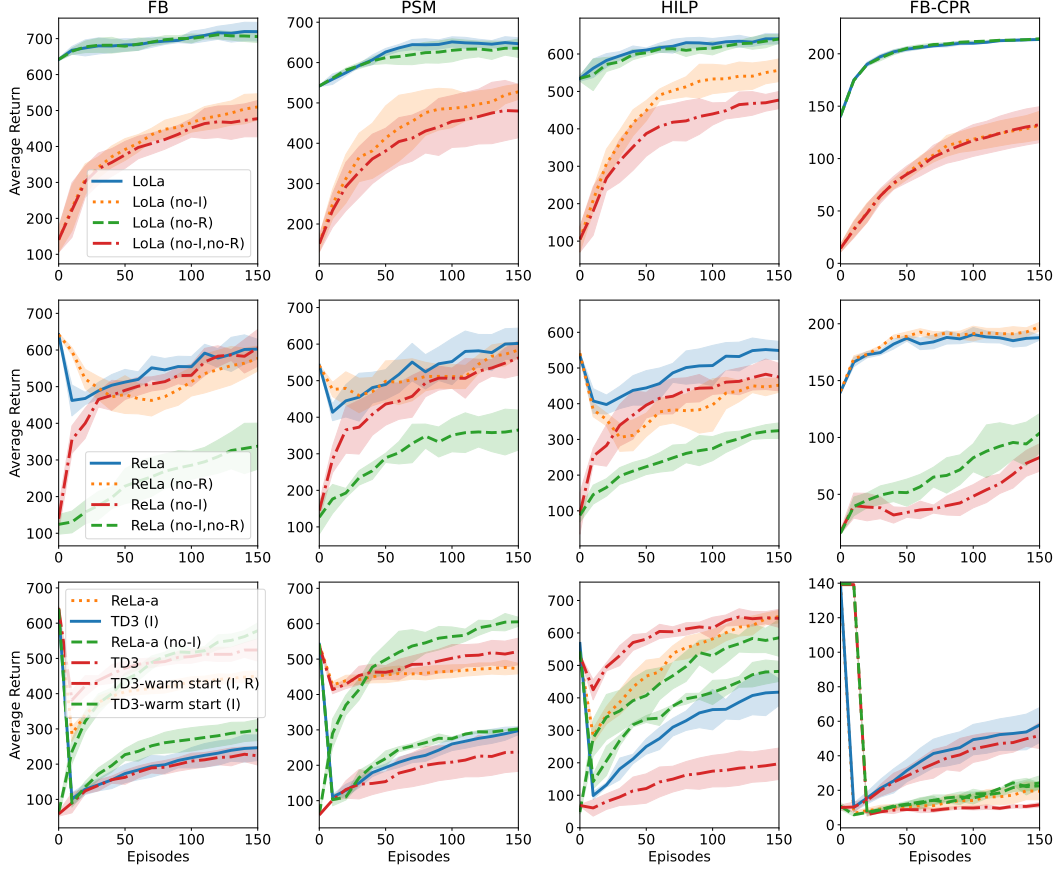


Figure 4: Average returns for several variations of LoLA, ReLA, and action-based TD3 with warm start. We use `no-R` to denote that we do not use the BFM’s estimated value function (i.e., for LoLA we do not bootstrap the terminal state and for ReLA we learn a critic from scratch) and `no-I` to denote that we do not use zero-shot policy initialization. Finally, for TD3 we use `R` to denote that we use residual critic since the standard implementation learns a critic from scratch.

299 (we use 0.98 and a lookahead of 100 or 250); this combination significantly reduces the role of
 300 the bootstrapped Q-function (discounted by 0.13 or 0.006). When looking at ReLA, residual critic
 301 helps in DMC domains but not in the HumEnv, where zero-shot initialization is the most important
 302 dimension. On the other hand, when zero-shot initialization is disabled in the DMC domains, the
 303 importance of residual critic is particularly evident and leads to almost match the performance of the
 304 best algorithm. Finally, the residual critic is also very important when performing direct adaptation
 305 in the action-space and helps in a faster recovery from the initial drop.

306 4.3 Dissecting the Suboptimality of Zero-Shot RL Policies

307 The previous results show that BFM’s are indeed learning skills that contain good policies for all the
 308 downstream tasks we study. This then raises the question on what causes the suboptimality of the
 309 zero-shot policy and the need of performing online adaptation to actually recover a better policy.
 310 We run a series of ablations with FB, HILP and PSM. We do not consider FB-CPR because our
 311 ablation involves offline training and we do not have access to an offline dataset for this model since
 312 it was trained online. We consider TD3 as learning algorithm since it is the building block of all the
 313 three BFM’s and focus on latent policy adaptation (we call this approach TD3-z to avoid confusion
 314 with TD3 used in the previous sections to optimize the full policy network). For all experiments in
 315 this section, we consider the standard scenario of **training from scratch**, no zero-shot initialization
 316 and no-residual critic. Specifically when searching in z space, we use a pretrained BFM actor and




Figure 5: Performance improvement w.r.t. the zero-shot policy for a TD3-based method trained from scratch for 3M steps to perform search in the latent policy space (i.e., TD3-z). We report the results for both online and offline training using the ExoRL (Yarats et al., 2022) dataset. We also ablate learning with the true task reward and the reward reconstructed by the BFM methods. We average the results over all the task of the Walker, Quadruped and Cheetah domains. We report the average performance over 5 seeds. We additionally show the performance of ReLA after training for 200 episodes.

317 initialize z along with the critic randomly and when learning in action space we initialize both
 318 the actor and critic randomly. We report the performance of TD3-z after 3M training steps when
 319 using the true reward function and the reward function reconstructed by the BFMs⁶, both offline and
 320 online. We do not consider `pointmass` in this test since TD3 does not work well when trained
 321 online due to the challenging exploration in the long-horizon tasks considered in this domain.

322 Overall, these experiments confirm that BFMs can express much better policies than the zero-shot
 323 policy and that optimizing for true rewards is crucial to unlock their full performance. When opti-
 324 mizing for the latent reward, offline TD3-z can already improve the zero-shot performance revealing
 325 the difficulty of optimizing all policies $\{\pi_z\}$ simultaneously during the pre-training process.⁷ Inter-
 326 estingly, when moving to online training on the latent reward performance can even drop. We con-
 327 jecture the cause is the distribution shift between online and offline samples. Given that the models
 328 were trained offline, their reward prediction degrades on out-of-distribution samples encountered
 329 during online adaptation, further skewing towards learning policies that are even less correlated to
 330 the true reward. This is confirmed when looking at the performance when optimizing for the true
 331 reward, which consistently lead to better results across offline and online tests, with online methods
 332 being overall better. This ablation confirms that focusing on searching in the z -space, while correct-
 333 ing the embedding errors is the right strategy to achieve fast adaptation online. Indeed, we see that
 334 ReLA recovers better policies than the one obtained by training from scratch TD3-z online for 3M
 335 episodes in only 200 episodes. Even faster if we use LoLA. This shows that leveraging information
 336 from the BFM is useful in many cases.

337 5 Related Work

338 **Unsupervised RL pre-training:** For language and vision, unsupervised pretraining has paved the
 339 way to extracting meaningful structure from data, scaling up, and obtaining impressive results for
 340 transfer and zero-shot generalization to different downstream tasks. In recent years, approaches have
 341 been proposed for unsupervised reinforcement learning: a training paradigm where a learning agent
 342 attempts to extract world structure and representations that will later allow it to solve diverse multi-

⁶Latent or reconstructed reward is given by $\tilde{r}_z(s) = \varphi(s) \cdot z_r$

⁷The fact that performance of HILP and PSM improve by about 10-15% by offline training on the reconstructed reward might be due to a non-perfect pre-train. Indeed, the pre-training condition should ensure that the actor is already optimal on any reconstructed reward on the training data distribution.

step decision-making problems in the environment. Various objectives have been proposed: world modeling (Bruce et al., 2024; Hansen et al., 2023), intrinsic rewards (Schmidhuber, 2019; Stadie et al., 2015; Sekar et al., 2020; Pathak et al., 2017), empowerment and mutual information skill learning (Klyubin et al., 2005; Eysenbach et al., 2018; Rajeswar et al., 2023; Gregor et al., 2016), goal-reaching (Ma et al., 2022; Park et al., 2023; 2024), successor measures (Touati & Ollivier, 2021; Agarwal et al., 2024; Sikchi et al., 2024) among others. In this work, we restrict our focus to the class of unsupervised RL objectives that learn a family of policies and allow us to query for a near-optimal policy given any test-time reward function without further learning or planning in the environment. Approaches belonging to this class often learn a state representation and use that to define the class of reward functions for which they learn the set of optimal policies. At inference time, they output the policy that is optimal for the projection of the reward function to this class of reward functions.

Fine-tuning and adaptation with unsupervised RL models: Similar to supervised pre-training approaches, unsupervised zero-shot RL models are not expected to output optimal policies for the given task. Rather they are expected to output a reasonable policy initialization that can be later fine-tuned or adapted. Prior approaches for policy adaptation with pre-trained RL models have mostly studied the offline-to-online setting. In this setting, a policy is trained with reward-labeled transitions first using offline data with specialized offline RL algorithms and then allowed to fine-tune by interacting with the environment and *retaining access to the offline data*. Offline RL algorithms (Levine et al., 2020; Sikchi et al., 2023) incorporate pessimism to avoid overestimation by restricting the policy to visit states closer to the dataset. Naively using the same algorithm to finetune online has been observed to lead to slow performance improvements and using an online RL algorithm leads to performance collapse at the beginning of fine-tuning (Luo et al., 2023). This behavior has been attributed to a distribution shift for critic-learning (Yu & Zhang, 2023), and prior works have investigated various techniques, such as calibration of Q-functions to mitigate this problem (Nakamoto et al., 2023). Our work, considers a different but practical paradigm for adaptation where a) no offline data is retained during finetuning and b) we learn from reward-free transitions. First, by only retaining pre-trained models we reduce compute requirements of learning from large pre-training offline datasets (Zhou et al., 2024), and by considering reward-free transitions we have a single model that can adapt to any downstream task (Kim et al., 2024). Learning online without retaining offline data suffers a significant initial drop of performance with respect to the pre-trained policy as recently investigated by Zhou et al. (2024).

6 Conclusion

Unsupervised zero-shot RL pre-training can result in an agent (a type of Behavioral Foundation Model, BFM) capable of accomplishing a wide variety of tasks having a noticeable but expected degree of suboptimality. This paper investigates and addresses the question of how to adapt these agents to be better at a task specified during test-time with limited environment interactions. We propose two fast adaptation strategies (LoLA and ReLA); The key insight behind our methods is to reuse pre-trained knowledge from BFM strategically and search over the learned latent policy space that provides a low-dimensional landscape favorable for gradient-based optimization. We have demonstrated the effectiveness of these strategies across various zero-shot BFMs. Notably, LoLA, an actor-only adaptation algorithm, demonstrates monotonic performance improvement on all domains and BFMs, making it a reliable choice when privileged resets are permitted. However, our findings also reveal an initial performance drop when employing any actor-critic method, including our proposed ReLA algorithm. This highlights the need for further investigation into mitigating forgetting in the actor-critic class of approaches. Future research directions include exploring meta-learning adaptation techniques, including in-context adaptation by learning to adapt in multi-task settings to optimize learning costs and improve overall performance.

391 **Broader Impact Statement**

392 Our work seeks to advance the adaptability of learning agents that interact with the environment. Our
 393 work pushes the frontier on scalable and adaptable agents by building upon unsupervised learning
 394 objectives that allow us to reuse a single trained model capable of a variety of tasks and propos-
 395 ing approaches that make these models more proficient at a task specified during test time rapidly
 396 without retaining any data used for pretraining. Prior works have studied a variety of objectives for
 397 unsupervised RL but fall short in their investigation of adaptation or propose a strategy that unlearns
 398 the policy before starting to improve performance. Our proposed approach is the first to our knowl-
 399 edge, to demonstrate monotonic performance improvement for pretrained RL agents without access
 400 to training data. There are potential societal consequences of our work, none which we feel must be
 401 specifically highlighted here.

402 **References**

- 403 Siddhant Agarwal, Harshit Sikchi, Peter Stone, and Amy Zhang. Proto successor measure:
 404 Representing the space of all possible solutions of reinforcement learning. *arXiv preprint*
 405 *arXiv:2411.19418*, 2024.
- 406 Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin,
 407 Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learn-
 408 ing from human feedback in llms. *arXiv preprint arXiv:2402.14740*, 2024.
- 409 André Barreto, Will Dabney, Rémi Munos, Jonathan J Hunt, Tom Schaul, Hado P van Hasselt,
 410 and David Silver. Successor features for transfer in reinforcement learning. *Advances in neural*
 411 *information processing systems*, 30, 2017.
- 412 Léonard Blier, Corentin Tallec, and Yann Ollivier. Learning successor states and goal-dependent
 413 values: A mathematical viewpoint. *arXiv preprint arXiv:2101.07123*, 2021.
- 414 Diana Borsa, André Barreto, John Quan, Daniel Mankowitz, Rémi Munos, Hado van Hasselt,
 415 David Silver, and Tom Schaul. Universal successor features approximators. *arXiv preprint*
 416 *arXiv:1812.07626*, 2018.
- 417 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
 418 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
 419 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- 420 Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes,
 421 Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative inter-
 422 active environments. In *Forty-first International Conference on Machine Learning*, 2024.
- 423 Yuri Burda, Harrison Edwards, Amos J. Storkey, and Oleg Klimov. Exploration by random network
 424 distillation. In *7th International Conference on Learning Representations, ICLR 2019, New Or-*
 425 *leans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL [https://openreview.net/](https://openreview.net/forum?id=H1lJJnR5Ym)
 426 [forum?id=H1lJJnR5Ym](https://openreview.net/forum?id=H1lJJnR5Ym).
- 427 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and
 428 Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of*
 429 *the IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.
- 430 Edoardo Cetin, Ahmed Touati, and Yann Ollivier. Finer behavioral foundation models via auto-
 431 regressive features and advantage weighting. *arXiv preprint arXiv:2412.04368*, 2024.
- 432 Peter Dayan. Improving generalization for temporal difference learning: The successor representa-
 433 tion. *Neural computation*, 5(4):613–624, 1993.

- 434 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
435 bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of*
436 *the North American chapter of the association for computational linguistics: human language*
437 *technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.
- 438 Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Ayzaan Wahid,
439 Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, et al. Palm-e: An embodied mul-
440 timodal language model. 2023.
- 441 Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need:
442 Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- 443 Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-
444 critic methods. In *International conference on machine learning*, pp. 1587–1596. PMLR, 2018.
- 445 Karol Gregor, Danilo Jimenez Rezende, and Daan Wierstra. Variational intrinsic control. *arXiv*
446 *preprint arXiv:1611.07507*, 2016.
- 447 Nicklas Hansen, Hao Su, and Xiaolong Wang. Td-mpc2: Scalable, robust world models for contin-
448 uous control. *arXiv preprint arXiv:2310.16828*, 2023.
- 449 Steven Hansen, Will Dabney, Andre Barreto, Tom Van de Wiele, David Warde-Farley, and
450 Volodymyr Mnih. Fast task inference with variational intrinsic successor features. *arXiv preprint*
451 *arXiv:1906.05030*, 2019.
- 452 Tobias Johannink, Shikhar Bahl, Ashvin Nair, Jianlan Luo, Avinash Kumar, Matthias Loskyll,
453 Juan Aparicio Ojea, Eugen Solowjow, and Sergey Levine. Residual reinforcement learning for
454 robot control. In *2019 international conference on robotics and automation (ICRA)*, pp. 6023–
455 6029. IEEE, 2019.
- 456 Junsu Kim, Seohong Park, and Sergey Levine. Unsupervised-to-online reinforcement learning.
457 *arXiv preprint arXiv:2408.14785*, 2024.
- 458 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International*
459 *Conference on Learning Representations (ICLR)*, 2015.
- 460 Alexander S Klyubin, Daniel Polani, and Chrystopher L Nehaniv. Empowerment: A universal
461 agent-centric measure of control. In *2005 ieee congress on evolutionary computation*, volume 1,
462 pp. 128–135. IEEE, 2005.
- 463 Wouter Kool, Herke van Hoof, and Max Welling. Buy 4 reinforce samples, get a baseline for free!
464 2019.
- 465 Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tuto-
466 rial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- 467 Hao Liu and Pieter Abbeel. Aps: Active pretraining with successor features. In *International*
468 *Conference on Machine Learning*, pp. 6736–6747. PMLR, 2021.
- 469 Yicheng Luo, Jackie Kay, Edward Grefenstette, and Marc Peter Deisenroth. Finetuning from
470 offline reinforcement learning: Challenges, trade-offs and practical solutions. *arXiv preprint*
471 *arXiv:2303.17396*, 2023.
- 472 Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy
473 Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training.
474 *arXiv preprint arXiv:2210.00030*, 2022.
- 475 Zakaria Mhammedi, Dylan J. Foster, and Alexander Rakhlin. The power of resets in online rein-
476 forcement learning. In *NeurIPS*, 2024.

- 477 Ashvin Nair, Abhishek Gupta, Murtaza Dalal, and Sergey Levine. Awac: Accelerating online rein-
478 forcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.
- 479 Mitsuhiro Nakamoto, Yuexiang Zhai, Anikait Singh, Max Sobol Mark, Yi Ma, Chelsea Finn, Aviral
480 Kumar, and Sergey Levine. Cal-ql: Calibrated offline rl pre-training for efficient online fine-
481 tuning. *arXiv preprint arXiv:2303.05479*, 2023.
- 482 Seohong Park, Oleh Rybkin, and Sergey Levine. Metra: Scalable unsupervised rl with metric-aware
483 abstraction. *arXiv preprint arXiv:2310.08887*, 2023.
- 484 Seohong Park, Tobias Kreiman, and Sergey Levine. Foundation policies with hilbert representations.
485 *arXiv preprint arXiv:2402.15567*, 2024.
- 486 Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration
487 by self-supervised prediction. In *International conference on machine learning*, pp. 2778–2787.
488 PMLR, 2017.
- 489 Matteo Pirota, Andrea Tirinzoni, Ahmed Touati, Alessandro Lazaric, and Yann Ollivier. Fast im-
490 itation via behavior foundation models. In *The Twelfth International Conference on Learning*
491 *Representations*, 2024. URL <https://openreview.net/forum?id=qnWtw310jb>.
- 492 Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John
493 Wiley & Sons, 2014.
- 494 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language
495 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 496 Sai Rajeswar, Pietro Mazzaglia, Tim Verbelen, Alexandre Piché, Bart Dhoedt, Aaron Courville, and
497 Alexandre Lacoste. Mastering the unsupervised reinforcement learning benchmark from pixels.
498 In *International Conference on Machine Learning*, pp. 28598–28617. PMLR, 2023.
- 499 Krishan Rana, Ming Xu, Brendan Tidd, Michael Milford, and Niko Sünderhauf. Residual skill
500 policies: Learning an adaptable skill-based action space for reinforcement learning for robotics.
501 In *Conference on Robot Learning*, pp. 2095–2104. PMLR, 2023.
- 502 Juergen Schmidhuber. Reinforcement learning upside down: Don’t predict rewards—just map them
503 to actions. *arXiv preprint arXiv:1912.02875*, 2019.
- 504 Ramanan Sekar, Oleh Rybkin, Kostas Daniilidis, Pieter Abbeel, Danijar Hafner, and Deepak Pathak.
505 Planning to explore via self-supervised world models. In *International conference on machine*
506 *learning*, pp. 8583–8592. PMLR, 2020.
- 507 Harshit Sikchi, Qinqing Zheng, Amy Zhang, and Scott Niekum. Dual rl: Unification and new
508 methods for reinforcement and imitation learning, 2023.
- 509 Harshit Sikchi, Siddhant Agarwal, Pranaya Jajoo, Samyak Parajuli, Caleb Chuck, Max Rudolph,
510 Peter Stone, Amy Zhang, and Scott Niekum. Rl zero: Zero-shot language to behaviors without
511 any supervision. *arXiv preprint arXiv:2412.05718*, 2024.
- 512 Tom Silver, Kelsey Allen, Josh Tenenbaum, and Leslie Kaelbling. Residual policy learning. *arXiv*
513 *preprint arXiv:1812.06298*, 2018.
- 514 Bradley C Stadie, Sergey Levine, and Pieter Abbeel. Incentivizing exploration in reinforcement
515 learning with deep predictive models. *arXiv preprint arXiv:1507.00814*, 2015.
- 516 Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- 517 Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Bud-
518 den, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy P. Lillicrap, and Martin A.
519 Riedmiller. Deepmind control suite. *CoRR*, abs/1801.00690, 2018.

- 520 Andrea Tirinzoni, Ahmed Touati, Jesse Farebrother, Mateusz Guzek, Anssi Kanervisto, Yingchen
521 Xu, Alessandro Lazaric, and Matteo Pirota. Zero-shot whole-body humanoid control via behav-
522 ior foundation models. In *The Thirteenth International Conference on Learning Representa-*
523 *tions*, 2025. URL <https://openreview.net/forum?id=9sOR0nYLTz>.
- 524 Ahmed Touati and Yann Ollivier. Learning one representation to optimize all rewards. In *NeurIPS*,
525 pp. 13–23, 2021.
- 526 Ahmed Touati, Jérémy Rapin, and Yann Ollivier. Does zero-shot reinforcement learning exist? In
527 *ICLR*. OpenReview.net, 2023.
- 528 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
529 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and
530 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 531 Yifan Wu, George Tucker, and Ofir Nachum. The laplacian in rl: Learning representations with
532 efficient approximations. *arXiv preprint arXiv:1810.04586*, 2018.
- 533 Denis Yarats, David Brandfonbrener, Hao Liu, Michael Laskin, Pieter Abbeel, Alessandro Lazaric,
534 and Lerrel Pinto. Don’t change the algorithm, change the data: Exploratory data for offline
535 reinforcement learning. *arXiv preprint arXiv:2201.13425*, 2022.
- 536 Zishun Yu and Xinhua Zhang. Actor-critic alignment for offline-to-online reinforcement learning.
537 In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and
538 Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*,
539 volume 202 of *Proceedings of Machine Learning Research*, pp. 40452–40474. PMLR, 23–29 Jul
540 2023. URL <https://proceedings.mlr.press/v202/yu23k.html>.
- 541 Tom Zahavy, Yannick Schroecker, Feryal Behbahani, Kate Baumli, Sebastian Flennerhag, Shaobo
542 Hou, and Satinder Singh. Discovering policies with domino: Diversity optimization maintaining
543 near optimality. *arXiv preprint arXiv:2205.13521*, 2022.
- 544 Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language
545 image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*,
546 pp. 11975–11986, 2023.
- 547 Zhiyuan Zhou, Andy Peng, Qiyang Li, Sergey Levine, and Aviral Kumar. Efficient online reinforce-
548 ment learning fine-tuning need not retain offline data. *arXiv preprint arXiv:2412.07762*, 2024.

Supplementary Materials

The following content was not necessarily subject to peer review.

6.1 Motivation for residual learning

Notation: We use matrix form and we identify for any $z \in \mathcal{Z}$, $\psi_z : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ and $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ by $\psi_z \in \mathbb{R}^{d \times |\mathcal{S} \times \mathcal{A}|}$ and $\phi \in \mathbb{R}^{d \times |\mathcal{S}|}$ respectively and $D_\rho = \text{diag}((\rho(s))_{s \in \mathcal{S}})$. Similarly, we identify for any $z \in \mathcal{Z}$, $F_z : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ and $B : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ by $F_z \in \mathbb{R}^{d \times |\mathcal{S} \times \mathcal{A}|}$ and $B \in \mathbb{R}^{d \times |\mathcal{S}|}$ respectively. We denote by $\Pi_\phi = \phi^\top (\phi D_\rho \phi^\top)^{-1} \phi D_\rho$ the $L^2(\rho)$ orthogonal projection onto the linear span of ϕ .

Proposition 1. Let $\phi : \mathcal{S} \rightarrow \mathbb{R}^d$ a state feature map and $\{\psi_z\}_{z \in \mathcal{Z}}$ the corresponding universal successor features for the policy family $\{\pi_z\}_{z \in \mathcal{Z}}$, i.e. $\psi_z(s, a) = \mathbb{E}[\sum_{t \geq 0} \gamma^t \phi(s_{t+1}) \mid (s, a), \pi_z]$. Then, for any reward function $r : \mathcal{S} \rightarrow \mathbb{R}$, we have: $Q_r^{\pi_z} = \psi_z^\top z_r + Q_{r - \phi^\top z_r}^{\pi_z}$ where $z_r = \mathbb{E}_{s \sim \rho}[\phi(s) \phi(s)^\top]^{-1} \mathbb{E}_{s \sim \rho}[\phi(s) r(s)]$, and $Q_{r - \phi^\top z_r}^{\pi_z}$ is the Q -function of the residual reward $r - \phi^\top z_r = (I - \Pi_\phi)r$.

Proof. for any reward function $r \in \mathbb{R}^{\mathcal{S}}$, we have $r = (\Pi_B + I - \Pi_B)r$, then

$$\begin{aligned} Q_r^{\pi_z} &= M^{\pi_z} r \\ &= M^{\pi_z} (\Pi_B + I - \Pi_B) r \\ &= M^{\pi_z} \Pi_B r + M^{\pi_z} (I - \Pi_B) r \\ &= M^{\pi_z} \phi^\top (\phi D_\rho \phi^\top)^{-1} \phi D_\rho r + M^{\pi_z} (r - \phi^\top (\phi D_\rho \phi^\top)^{-1} \phi D_\rho r) \\ &= \psi_z^\top z_r + Q_{r - \phi^\top z_r}^{\pi_z} \end{aligned}$$

where the last equation follows from the fact that $\psi_z^\top = M^{\pi_z} \phi^\top$ and $z_r = \mathbb{E}_{s \sim \rho}[\phi(s) \phi(s)^\top]^{-1} \mathbb{E}_{s \sim \rho}[\phi(s) r(s)] = (\phi D_\rho \phi^\top)^{-1} \phi D_\rho r$ \square

Proposition 2. Let assume that for any $z \in \mathcal{Z}$, F_z is a stationary point of the FB training loss $\ell(F, B)$, namely, the functional derivative $\frac{\partial \ell}{\partial F_z}$ of the loss with respect of F_z is 0. Then,

$$\begin{aligned} Q_r^{\pi_z} &= F_z^\top z_r + Q_{(I - \Pi_B)r}^{\pi_z} \\ &= F_z^\top z_r + (M^{\pi_z} - F_z^\top B D_\rho) (r - \Pi_B r) \end{aligned}$$

where $z_r = \mathbb{E}_{s \sim \rho}[B(s) r(s)]$.

Proof. Let's remind the FB training loss:

$$\ell(F, B) = \mathbb{E}_{z, (s, a) \sim \rho} \left[\left(F(s, a, z)^\top B(s^+) - P(ds^+ \mid s, a) / \rho(ds^+) - (P^{\pi_z} \bar{F})(s, a, z)^\top \bar{B}(s^+) \right)^2 \right]$$

In matrix form, we obtain:

$$\ell(F, B) = \mathbb{E}_z \left[\text{Trace} \left((F_z^\top B - P D_\rho^{-1} - \gamma P^{\pi_z} \bar{F}_z^\top \bar{B})^\top D_\rho (F_z^\top B - P D_\rho^{-1} - \gamma P^{\pi_z} \bar{F}_z^\top \bar{B}) D_\rho \right) \right]$$

571 if F_z satisfies the stationarity conditions, i.e., $\frac{\partial \ell}{\partial F_z} = 0$, then, we have

$$\begin{aligned}
\frac{\partial \ell}{\partial F_z} = 0 &\Rightarrow 2BD_\rho (F_z^\top B - PD_\rho^{-1} - \gamma P^{\pi_z} F_z^\top B)^\top D_\rho = 0 \\
&\Rightarrow 2D_\rho (F_z^\top B - PD_\rho^{-1} - \gamma P^{\pi_z} F_z^\top B) D_\rho B^\top = 0 \\
&\Rightarrow F_z^\top BD_\rho B^\top = PB^\top + \gamma P^{\pi_z} F_z^\top BD_\rho B^\top \\
&\Rightarrow F_z^\top = M^{\pi_z} B^\top (BD_\rho B^\top)^{-1} \\
&\Rightarrow F_z^\top BD_\rho = M^{\pi_z} B^\top (BD_\rho B^\top)^{-1} \\
&\Rightarrow F_z^\top BD_\rho = M^{\pi_z} B^\top (BD_\rho B^\top)^{-1} BD_\rho
\end{aligned}$$

572 Therefore $F_z^\top BD_\rho = M^{\pi_z} \Pi_B$ where $\Pi_B = B^\top (BD_\rho B^\top)^{-1} BD_\rho$ is the $L^2(\rho)$ orthogonal pro-
573 jection onto the linear span of B .

574 Let Π_{B^\perp} the orthogonal projection onto the orthogonal of B . By definition, we have $\Pi_{B^\perp} = I - \Pi_B$

575 We have:

$$\begin{aligned}
Q_r^{\pi_z} &= M^{\pi_z} r \\
&= M^{\pi_z} (\Pi_B + \Pi_{B^\perp}) r \\
&= M^{\pi_z} \Pi_B r + M^{\pi_z} \Pi_{B^\perp} r \\
&= F_z^\top BD_\rho r + M^{\pi_z} \Pi_{B^\perp} r \\
&= F_z^\top z_r + M^{\pi_z} \Pi_{B^\perp} r
\end{aligned}$$

576 where $z_r = BD_\rho r = \mathbb{E}_{s \sim \rho} [B(s)r(s)]$

577 Therefore, we have:

$$Q_r^{\pi_z} = F_z^\top z_r + Q_{\Pi_{B^\perp} r}^{\pi_z}$$

578 where the second term is the the Q-function of the residual reward $\Pi_{B^\perp} r$

579 Moreover, since $\Pi_{B^\perp}^2 = \Pi_{B^\perp}$, we can write:

$$\begin{aligned}
Q_r^{\pi_z} &= F_z^\top z_r + (M^{\pi_z} \Pi_{B^\perp}) (\Pi_{B^\perp} r) \\
&= F_z^\top z_r + (M^{\pi_z} - F_z^\top BD_\rho) (r - \Pi_B r)
\end{aligned}$$

580 Which means that the residual term can be expressed as the successor measure approximation error
581 (due to the low-rank decomposition of FB model) multiplied by the reward error (due to the reward
582 embedding in the span of B). \square

583 7 Pseudocode

584 Algorithm 1 and 2 outline pseudocode for ReLA and LoLA respectively.

585 8 Experimental Setup

586 8.1 Environments

587 We list the continuous control environments from the DeepMind Control Suite (Tassa et al., 2018)
588 and Humenv (Tirinzoni et al., 2025) used in this work in Table 2.

589 8.2 Behavioral Foundation Models

590 We trained all the BFMs except for the FB-CPR model that is publicly available (code link).

Algorithm 1: ReLA

Load Frozen BFM’s successor features $\psi(s, a, z)$ and policy $\pi_z(s)$ networks.

Initialize residual critic networks $Q_{\theta_1}^{\text{residual}}, Q_{\theta_2}^{\text{residual}}$, replay buffer $\mathcal{D}_{\text{online}}$, exploration std σ . Update to Data ratio (UTD) M , Initialize target networks: $Q_{\theta_1'}^{\text{residual}} \leftarrow Q_{\theta_1}^{\text{residual}}, Q_{\theta_2'}^{\text{residual}} \leftarrow Q_{\theta_2}^{\text{residual}}$.

Compute zero-shot latent z_r using inference samples for the BFM agent with test-time reward function.

for each environment step t **do**

Select $a_t = \pi_z(s_t) + \epsilon, \epsilon \sim \mathcal{N}(0, \sigma)$

Execute a_t ; observe r_t, s_{t+1}

Store (s_t, a_t, r_t, s_{t+1}) in $\mathcal{D}_{\text{online}}$

Sample M mini-batches

Batch $_i = \{(s_i, a_i, r_i, s'_i)\} \sim \mathcal{D}_{\text{online}}$

Compute target Q-value:

$y_i = r_i + \gamma(\psi(s'_i, \pi_z(s'_i), z_r) \cdot z_r + \min\{Q_{\theta_1'}(s'_i, \pi_z(s'_i)), Q_{\theta_2'}(s'_i, \pi_z(s'_i))\})$

TemporalDifferenceUpdate($\psi(s_i, a_i, z_r) \cdot z_r + Q_{\theta_{k \in [1,2]}^{\text{residual}}}(s_i, y_i)$ for $i \in [M]$ using critic parameterization from Eq 5

Latent Policy Update Update z taking gradient step as in Eq 6 on $\cup_{i \in [m]} \text{Batch}_i$.

Update target networks by polyak averaging;

end

Algorithm 2: LoLA

Load Frozen BFM’s successor features $\psi(s, a, z)$ and policy $\pi_z(s)$ networks

Initialize latent policy $\pi_{\mu, \sigma} = \mathcal{N}(\mu = z_r, \sigma)$, replay buffer $\mathcal{D}_{\text{online}}$, sampling state distribution $\nu(\mathcal{D}_{\text{online}}, d_0)$, z budget k , initial state budget m , horizon n

Compute zero-shot latent z_r using inference samples for the BFM agent with test-time reward function.

for each gradient step **do**

for $b=1..m$ **do**

$s_0 \sim \mu(\mathcal{D}_{\text{online}}, d_0)$

for $i=1..k$ **do**

$z_b^i \sim \pi_{\mu, \sigma}$, Reset to s_0

Rollout trajectory τ_b^i by taking actions given by $a_t = \pi_{z_b^i}(s_t)$

Compute $R(s_0, z_b^i)$

Collect states from τ_b^i in $\mathcal{D}_{\text{online}}$

end

end

Update $\pi_{\mu, \sigma}$ by taking gradient step in Eq 7.

end

Figure 6: Pseudocode of our proposed adaptation methods: Residual Latent Adaptation (ReLA) and Lookahead Latent Adaptation (LoLA).

Domain	Observation dimension	Action dimension	Episode length
Pointmass	4	2	1000
Walker	24	6	1000
Cheetah	17	6	1000
Quadruped	78	12	1000
HumEnv	358	69	300

Table 2: Overview of observation spaces, action spaces and episode length of environments used in this work.

591 **Offline BFMs.** We train the BFMs using the publicly available dataset from ExoRL (Yarats et al.,
 592 2022) collected using the RND algorithm (Burda et al., 2019). We used the authors implementation
 593 for FB and FB-CPR (code link) and reimplemented PSM and HILP. We report in Table 3 the set of
 594 hyperparameter used for the algorithms.

595 **FB architecture.** The backward representation network $B(s)$ is represented by a feedforward neu-
 596 ral network with two hidden layers, each with 256 units, that takes as input a state and outputs a
 597 d -dimensional embedding. For the forward network $F(s, a, z)$, we first preprocess separately (s, a)
 598 and (s, z) by two feedforward networks with one single hidden layer (with 1024 units) to 512-
 599 dimensional space. Then we concatenate their two outputs and pass it into two heads of feedforward
 600 networks (each with one hidden layer of 1024 units) to output a d -dimensional vector. For the pol-
 601 icy network $\pi(s, z)$, we first preprocess separately s and (s, z) by two feedforward networks with
 602 one single hidden layer (with 1024 units) to 512-dimensional space. Then we concatenate their two
 603 outputs and pass it into another one single hidden layer feedforward network (with 1024 units) to
 604 output to output a d_A -dimensional vector, then we apply a Tanh activation as the action space is
 605 $[-1, 1]^{d_A}$.

606 For all the architectures, we apply a layer normalization and Tanh activation in the first layer in
 607 order to standardize the states and actions. We use ReLU for the rest of layers. We also pre-
 608 normalize $z : z \leftarrow \sqrt{d} \frac{z}{\|z\|_2}$ in the input of F , and π .

Table 3: BFM hyperparameters. We largely reuse the hyperparameters from [Pirrotta et al. \(2024\)](#) for FB, from [Park et al., 2024](#) for HILP.

	Hyperparameter	Walker	Cheetah	Quadruped	Pointmass	
FB	Forward Backward (Touati & Ollivier, 2021)	Embedding Dimension d	100	50	50	100
		Embedding Prior	S^d	S^d	S^d	S^d
		Embedding Prior Goal Prob.	0.5	0.5	0.5	0.5
		B Normalization	ℓ_2	ℓ_2	ℓ_2	ℓ_2
		Orthonormal Loss Coeff.	1	1	1	1
	Optimizer (Adam) (Kingma & Ba, 2015)	Learning Rate (F, B)	$(10^{-4}, 10^{-4})$	$(10^{-4}, 10^{-4})$	$(10^{-4}, 10^{-4})$	$(10^{-4}, 10^{-6})$
		Learning Rate (π)	10^{-4}	10^{-4}	10^{-4}	10^{-6}
		Target Network EMA	0.99	0.99	0.99	0.99
HILP	Hilbert Representations (Park et al., 2024)	Embedding Dimension d	50	50	50	100
		Feature Learning Expectile	0.5	0.5	0.5	0.5
		Feature Learning Discount Factor	0.98	0.98	0.98	0.98
		Successor feature loss	Q-loss	Q-loss	Q-loss	Q-loss
	Optimizer (Adam)	Learning Rate (SF, F)	$(10^{-4}, 10^{-5})$	$(10^{-4}, 10^{-4})$	$(10^{-4}, 10^{-4})$	$(10^{-4}, 10^{-4})$
		Learning Rate (π)	10^{-4}	10^{-4}	10^{-4}	10^{-4}
		Target Network EMA features	0.995	0.995	0.995	0.995
		Target Network EMA SF	0.99	0.99	0.99	0.99
PSM	Proto Successor Measures (Agarwal et al., 2024)	Embedding Dimension d	100	50	50	100
		Policy Codebook Size	2^{16}	2^{16}	2^{16}	2^{16}
		Feature Learning Timesteps	400k	400k	400k	400k
		Embedding Prior Goal Prob.	0.5	0.5	0.5	0.5
		B Normalization	ℓ_2	ℓ_2	ℓ_2	ℓ_2
		Orthonormal Loss Coeff.	1	1	1	1
	Optimizer (Adam) (Kingma & Ba, 2015)	Learning Rate (F, B)	$(10^{-4}, 10^{-4})$	$(10^{-4}, 10^{-4})$	$(10^{-4}, 10^{-4})$	$(10^{-4}, 10^{-6})$
		Learning Rate (π)	10^{-4}	10^{-4}	10^{-4}	10^{-6}
		Target Network EMA	0.99	0.99	0.99	0.99
	Policy (TD3) (Fujimoto et al., 2018)	Target Policy Noise	$\mathcal{N}(0, 0.2)$	$\mathcal{N}(0, 0.2)$	$\mathcal{N}(0, 0.2)$	$\mathcal{N}(0, 0.2)$
		Target Policy Clipping	0.3	0.3	0.3	0.3
		Policy Update Frequency	1	1	1	1
Common	Batch Size	1024	1024	1024	1024	
	Gradient Steps	3M	3M	3M	3M	
	Discount Factor γ	0.98	0.98	0.98	0.99	
	Reward Inference Samples	250,000	250,000	250,000	250,000	
	ExoRL number of trajectories	5,000	5,000	5,000	5,000	

HILP architecture. We use the same policy architecture as FB as well F-architecture for the successor features. We learn the HILP features using a 2 layers MLP with hidden dimension 1024. Even in this case, z is normalized.

PSM architecture. We use the same policy architecture as FB as well F-architecture for the successor features. We learn the PSM features using a 2 layers MLP with hidden dimension 256. Even in this case, z is normalized.

8.3 Algorithm Implementation

All the actor-critic algorithms are implemented using TD3 ([Fujimoto et al., 2018](#)) as the base off-policy algorithm. When learning from scratch we use a 2 two layers MLP with hidden dimension 1024 and ReLU activation. We use the same configuration also for the critic.

For ReLA, we use a small 2 layer MLPs with 64 hidden dimensions and ReLU activation as residual network. In the ablation, when residual critic is deactivated, we use the same critic network as for standard TD3.

For LoLA, we use a Gaussian policy centered around the learned z and learn simultaneously mean and standard deviation.

8.4 Hyperparameters

For all baselines and our method, we run a hyperparameter sweep across domains and tasks and choose the configuration that performs the best across tasks for each BFM.

TD3-based algorithms. We run a hyperparameter sweep on Update to Data ratio (UTD) in $[1, 4, 8]$, actor update in frequency in $[1, 4]$. We use a small 2 layer MLP with 64 hidden nodes for the residual network which we found to work best for fast adaptation. When not using residual critic, we learn a critic from scratch using a 2 layer MLP with 1024 hidden nodes. We use 10^{-4} as learning rate for both critic and actor. We use either warm start of 0 steps or 5000 steps.

LoLA. We consider hyperparameter sweep between a lookahead horizons of $[50, 100, 250]$, the number of total trajectories per update to be 10, and number of trajectories for a sampled state to be 5 (for calculating baseline). We sweep between $[0, 0.2, 0.5]$ for the probability of resetting to initial state distribution and otherwise sampling from states encountered in replay buffer. We sweep also the learning rate in $[0.1, 0.05]$.

9 Additional Experiments

As mentioned in the main paper, `pointmass` is the domain where actor-critic algorithms incurs a significant initial drop. Figure 7 shows the average performance improvement without the `pointmass` domain. As we can see, ReLA has still a initial drop but it is much more reduced compared to what reported in the main. Previous papers (e.g. [Pirotta et al., 2024](#)), noticed that a smaller learning rate helped in `pointmass`. In our experiments we kept the learning rate fixed at 10^{-4} for all the domains, it would be interesting to test different values.

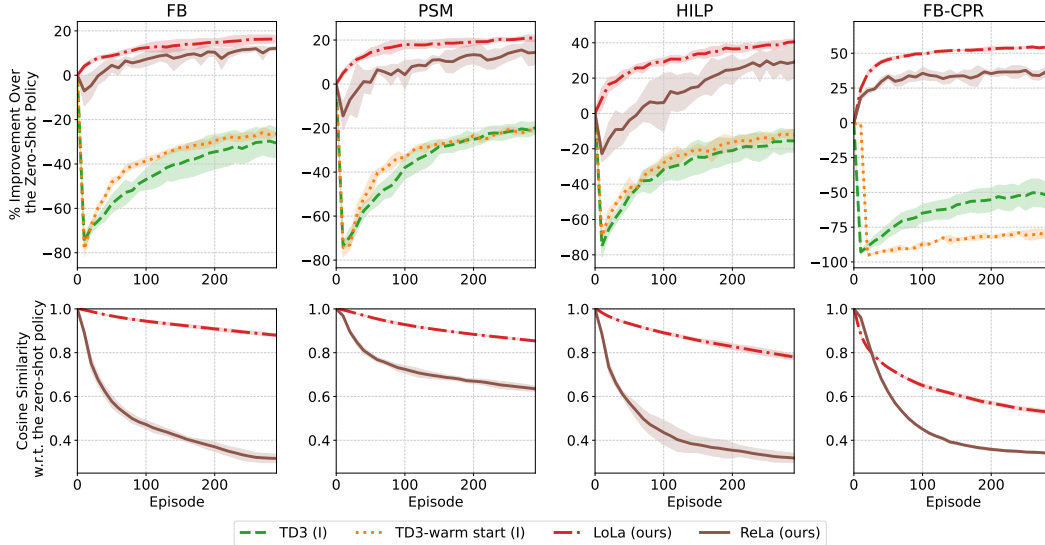


Figure 7: **Top:** Performance improvement w.r.t. the zero-shot policy for different online fast adaptation methods and BFMs without the `Pointmass` domain.

9.1 Per Algorithm Per Domain Ablation Studies

We conduct extensive ablation studies to understand the impact of key design choices in our methods, specifically: (1) zero-shot initialization in LoLA and ReLA variants, (2) value function bootstrapping in LoLA, (3) residual critics in ReLA variants and action-based TD3 with warm start. Table 4 provides a comprehensive list of the algorithm variants considered.

Algorithm	Zero-Shot Policy Initialization	Residual Critic ([†]) or Bootstrapped Return (+)	Critic Trained from scratch	Search space	WSRL
LoLA	✓	✓([†])		z	actor-only
LoLA (no-I)		✓([†])		z	
LoLA (no-R)	✓		✓	z	
LoLA (no-I, no-R)			✓	z	
ReLA	✓	✓([†])		z	actor-critic
ReLA-warm-start	✓	✓([†])		z	
ReLA (no-I)		✓([†])		z	
ReLA (no-R)	✓		✓	z	
ReLA-warm-start (no-R)	✓		✓	z	✓
ReLA (no-I, no-R)			✓	z	
ReLA-a	✓	✓([†])		a	actor-critic
ReLA-a (no-I)		✓([†])		a	
ReLA-a (no-R)	✓		✓	a	
ReLA-a (no-I, no-R)			✓	a	
TD3-z			✓	z	actor-critic
TD3 (I)	✓		✓	a	
TD3-warm-start (I) (i.e., using WSRL)	✓		✓	a	✓
TD3-warm-start (I, R)	✓	✓([†])		a	✓

Table 4: Summary of the algorithm variations considered in the main paper. Search space z means latent policy adaptation leveraging the policy space $\{\pi_z\}$ constructed by the BFM. Search space a denotes fine-tuning in action space (i.e., updating all policy network parameters).

We evaluated these variants across four DMC domains (Quadruped, Pointmass, Cheetah, Walker) using FB, HILP and PSM, and on HumEnv using FB-CPR, each experiment conducted over five random seeds. The results are shown in Figure 8, 9, 10 and 11.

Zero-Shot Initialization (no-zs-init): Removing zero-shot initialization consistently degraded early-stage performance across all methods and domains, with the only exception being ReLA-a with FB and PSM on pointmass. The benefit of zero-shot initialization is especially significant on LoLA.

Bootstrapping (no-bootstrap): We hypothesized that value function bootstrapping could help stabilizing LoLA. However, we did not notice such benefit from our ablation experiments.

Residual critics (no-residual): Removing residual critics in ReLA variants and TD3-based algorithm strongly impaired the effectiveness of the algorithm. This effect was especially pronounced for ReLA-a and TD3 on DMC domains.




Figure 8: Ablation studies evaluating adaptations of HILP on four DMC tasks (Quadruped, Pointmass, Cheetah, Walker). Experiments include disabling zero-shot initialization ("no-I") and/or removing residual critics ("no-R") from LoLA, ReLA, and three additional variants: (1) ReLA-a: update a instead of z in ReLA, (2) ReLA with warm start (ReLA-warm start), and (3) action-based TD3 with warm start (TD3-warm start). Shaded areas represent standard errors across 5 seeds.

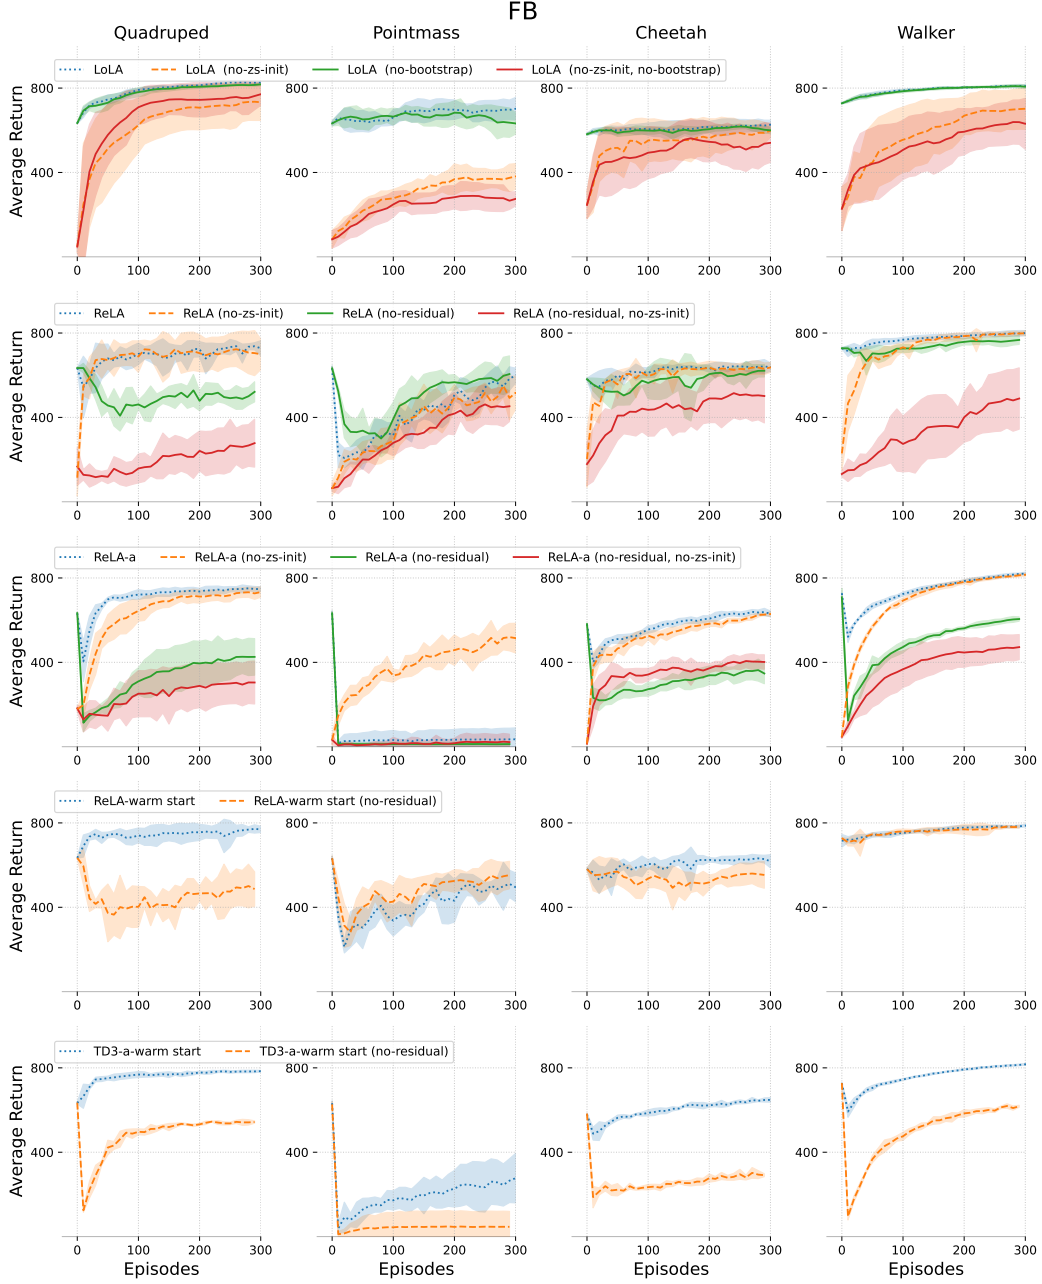


Figure 9: Ablation studies evaluating adaptations of FB on four DMC tasks (Quadruped, Point-mass, Cheetah, Walker). Experiments include disabling zero-shot initialization ("no-I") and/or removing residual critics ("no-R") from LoLA, ReLA, and three additional variants: (1) ReLA-a: update a instead of z in ReLA, (2) ReLA with warm start (ReLA-warm start), and (3) action-based TD3 with warm start (TD3-warm start). Shaded areas represent standard errors across 5 seeds.

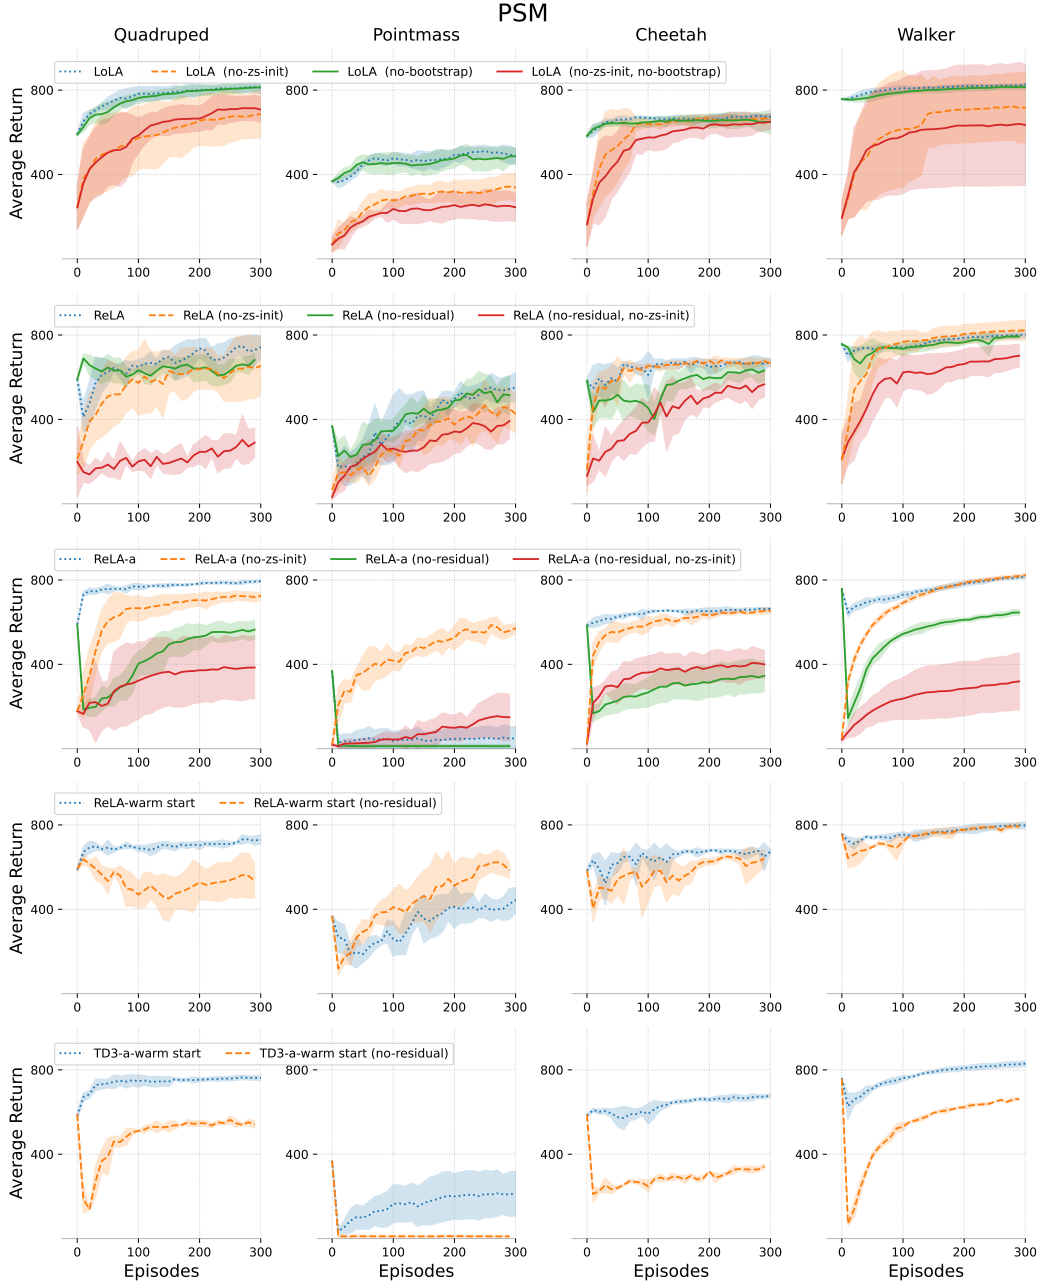


Figure 10: Ablation studies evaluating adaptations of PSM on four DMC tasks (Quadruped, Pointmass, Cheetah, Walker). Experiments include disabling zero-shot initialization ("no-I") and/or removing residual critics ("no-R") from LoLA, ReLA, and three additional variants: (1) ReLA-a: update action instead of z in ReLA, (2) ReLA with warm start (ReLA-warm start), and (3) action-based TD3 with warm start (TD3-warm start). Shaded areas represent standard errors across 5 seeds.

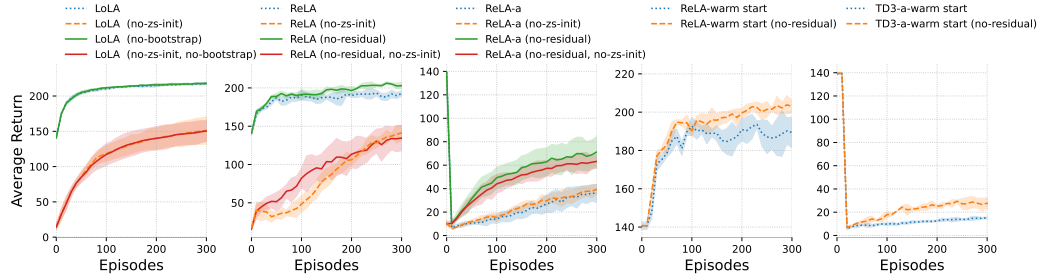


Figure 11: Ablation studies for adaptation with FB-CPR on 45 HumEnv tasks. Experiments include disabling zero-shot initialization ("no-I") and/or re-moving residual critics ("no-R") from LoLA, ReLA, and three additional variants: (1) ReLA-a: update action instead of z in ReLA, (2) ReLA with warm start (ReLA-warm start), and (3) action-based TD3 with warm start (TD3-warm start). Shaded areas represent standard errors across 5 seeds.