

Adapt On-the-Go: Behavior Modulation for Single-Life Robot Deployment

Annie S. Chen^{*1}, Govind Chada^{*1}, Laura Smith², Archit Sharma¹, Zipeng Fu¹, Sergey Levine², Chelsea Finn¹
Stanford University¹, UC Berkeley²

Abstract—To succeed in the real world, robots must cope with situations that differ from those seen during training. We study the problem of adapting on-the-fly to such novel scenarios during deployment, by drawing upon a diverse repertoire of previously-learned behaviors. Our approach, RObust Autonomous Modulation (ROAM), introduces a mechanism based on the perceived value of pre-trained behaviors to select and adapt pre-trained behaviors to the situation at hand. Crucially, this adaptation process all happens within a single episode at test time, without any human supervision. We provide theoretical analysis of our selection mechanism and demonstrate that ROAM enables a robot to adapt rapidly to changes in dynamics both in simulation and on a real Go1 quadruped, even successfully moving forward with roller skates on its feet. Our approach adapts over 2x as efficiently compared to existing methods when facing a variety of out-of-distribution situations during deployment by effectively choosing and adapting relevant behaviors on-the-fly.

I. INTRODUCTION

A major obstacle to the broad application of sequential decision-making agents is their inability to adapt to unexpected circumstances, which limits their uses largely to tightly controlled environments. Even equipped with prior experience and pre-training, agents will inevitably encounter out-of-distribution (OOD) situations at deployment time that may require a large amount of on-the-fly adaptation. In this work, we aim to enable a robot to autonomously handle novel scenarios encountered during deployment, while drawing upon a diverse set of pre-trained behaviors that may improve its versatility. We hypothesize that by doing some of the adaptation in the space of pre-trained behaviors (rather than only in the space of parameters), we can much more quickly react and adapt to novel circumstances. Accumulating a set of different behavior policies is a relatively straightforward task through online or offline episodic reinforcement learning, using different reward functions or skill discovery methods. However, the on-the-fly selection and adaptation of these behaviors during deployment, particularly in novel environments, present significant challenges. Consider tasking a quadrupedal robot that has acquired many different behaviors, e.g., walking, crouching, and limping via training in simulation, with a search-and-rescue mission in the real world. When deployed on this task with unstructured obstacles, the robot may bump into an obstacle and damage its leg, and it must be able to dynamically adapt its choice of behaviors to continue its mission with the damage.

Existing adaptation methods often operate within the standard, episodic training paradigm where the agent is assumed

to be reset and have another chance to attempt the task each time [1]–[3]. However, these are idealized conditions that in practice often rely on human intervention. In the above search-and-rescue example, the robot’s state cannot be arbitrarily restored; during deployment, the robot cannot be repaired, and it is not feasible for a human to fetch it if it gets stuck in a situation it is not equipped to handle. This situation necessitates adapting at test time both quickly and autonomously, to succeed at the task within a single episode. Therefore, we frame our problem setting as an instantiation of single-life deployment [4], where the agent is given prior knowledge but evaluated on its ability to successfully complete a given task later in a ‘single-life’ trial, during which there are no human-provided resets. The robot is provided with a diverse set of prior behaviors trained through episodic RL, and the single life poses a sequence of new situations.

To solve this problem, the robot must identify during deployment which behaviors are most suited to its situation at a given timestep and have the ability to fine-tune those behaviors in real time, as the pre-trained behaviors may not optimally accommodate new challenges. We introduce a simple method called RObust Autonomous Modulation (ROAM), which foremost aims to quickly identify the most appropriate behavior from its pre-trained set at each point in time during single-life deployment. Rather than introducing an additional component like a high-level controller to select behaviors, we leverage the value function of each behavior. The value function already estimates whether each policy will be successful, but may not be accurate for states that were not encountered during training of that behavior. Therefore, prior to deployment, we fine-tune each behavior’s value function with a regularized objective that encourages behavior identifiability: the regularizer is a behavior classification loss. Then, at each step during deployment, ROAM samples a behavior proportional to its classification probability, executes an action from that behavior, and optionally fine-tunes the selected behavior for additional adaptation.

The main contribution of this paper is a simple algorithm for autonomous, deployment-time adaptation to novel scenarios. At a given state, with the additional cross-entropy regularizer, ROAM can constrain each behavior’s value to be lower than the value of behaviors for which that state appears more frequently. Consequently, our method incentivizes each behaviors to differentiate between familiar and unfamiliar states, allowing ROAM to better recognize when a behavior will be useful. We conduct experiments on both simulated locomotion tasks and on a real Go1 quadruped robot. In simulation, our method completes the deployment task more than two times faster on

^{*} equal contribution. Correspondence to asc8@stanford.edu. See <https://anniesch.github.io/adapt-on-the-go/> for videos and full paper.

On-The-Go Adaptation by Efficiently Leveraging Prior Behaviors

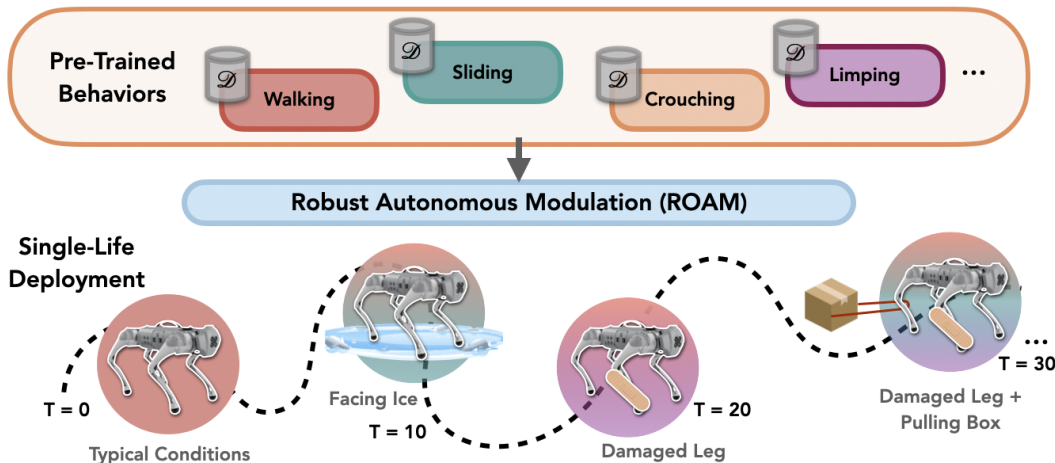


Fig. 1: **On-The-Go Adaptation via Robust Autonomous Modulation (ROAM)**. An agent will inevitably encounter a wide variety of situations during deployment, and handling such situations may require a variety of different behaviors. We propose Robust Autonomous Modulation (ROAM), which dynamically employs relevant behaviors as different situations arise during deployment.

average than existing methods, including two prior methods designed for fast adaptation. We also empirically analyze how the additional cross-entropy term in the loss function of ROAM contributes to more successful utilization of the prior behaviors. Furthermore, ROAM enables the Go1 robot to adapt on-the-go to various OOD situations without human interventions or supervision in the real world. With ROAM, the robot can successfully pull heavy luggage, pull loads with dynamic weights, and even slide forward with two roller skates on its front feet, even though it never encountered loads or wore roller skates during training.

II. RELATED WORK

We consider the problem of enabling an agent to act robustly when transferring to unstructured test-time conditions that are unknown at train-time. One common instantiation of this problem is transfer to different dynamics, e.g., in order to transfer policies trained in simulation to the real world. A popular approach in achieving transfer under dynamics shift is domain randomization, i.e., randomizing the dynamics during training [5]–[15] to learn a robust policy. Our approach is similar in that it takes advantage of different MDPs during training; however, a key component of ROAM is to leverage and modulate diverse skills rather than a single, robust policy. We find in Section V that challenging test-time scenarios may require distinct behaviors at different times, and we design our method to be robust to those heterogeneous conditions.

Another class of methods involve training policies that exhibit different behavior when conditioned on dynamics parameters, then distilling these policies into one that can be deployed in target domains where this information is not directly observable. The train-time supervision can come in the form of the parameter values [16], [17] or a learned representation of them [18], [19]. Thereafter, there are several ways prior work have explored utilizing online data to identify which behavior is appropriate on-the-fly, e.g., using search in latent space [11], [20], [21], or direct inference using proprioceptive history [18], [19], [22], or prediction based on egocentric depth [23]–[26]. In this work, we do not rely

on domain-specific information nor external supervision for when particular pre-trained behaviors are useful. Moreover, in contrast to many of the above works, we focus on solving tasks that may be OOD for all prior behaviors individually.

Meta-RL is another line of work that achieves rapid adaptation without privileged information by optimizing the adaptation procedure during training [2], [27]–[34] to be able to adapt quickly to a new situation at test-time. Meta-RL and the aforementioned domain randomization-based approaches entangle the training processes with the data collection, requiring a lot of *online* samples that are collected in a particular way for pre-training. The key conceptual difference in our approach is that ROAM is *agnostic to how the pre-trained policies and value functions are obtained*. Moreover, while meta-RL methods often use hundreds of pre-training tasks, or more, our approach can provide improvements in new situations even with a relatively small set of pre-trained behaviors (e.g. just 4 different behaviors improve performance in Section V). Other transfer learning approaches adapt the weights of the policy to a new environment or task, either through rapid zero-shot adaptation [4], [35], [36] or through extended episodic online training [37]–[41]. Unlike these works, we focus on adaptation within a single episode to a variety of different situations.

Another rich body of work considers how to combine prior behaviors to solve long-horizon tasks, and some of these works also focus on learning or discovering useful skills [42]–[49]. Many of these methods involve training a high-level policy that learns to compose learned skills into long-horizon behaviors [50]–[59]. We show in Section V that such a high-level policy may not be effective for on-the-go behavior selection. Moreover, our work does not focus on where the behaviors come from – they could be produced by these prior methods, or their rewards could be specified manually. Instead, we focus on quickly selecting and adapting the most suitable skill in an OOD scenario, without requiring an additional online training phase to learn a hierarchical controller.

III. PRELIMINARIES

In this section, we describe some preliminaries and formalize our problem statement. We are given a set of n prior behaviors, where each behavior i is trained through episodic RL for a particular MDP $\mathcal{M}_i = (\mathcal{S}, \mathcal{A}, \mathcal{P}_i, \mathcal{R}_i, \rho_0, \gamma)$ where \mathcal{S} is the state space, \mathcal{A} is the agent’s action space, $\mathcal{P}_i(s_{t+1}|s_t, a_t)$ represents the environment’s transition dynamics, $\mathcal{R}_i : \mathcal{S} \rightarrow \mathbb{R}$ indicates the reward function, $\rho_0 : \mathcal{S} \rightarrow \mathbb{R}$ denotes the initial state distribution, and $\gamma \in [0, 1)$ denotes the discount factor. Each of the n MDPs \mathcal{M}_i has potentially different dynamics and reward functions \mathcal{P}_i and \mathcal{R}_i , often leading to different state visitation distributions. Each behavior corresponds to a policy π_i and a value function V_i as well as a buffer of trajectories $\tau \in \mathcal{D}_i$ collected during this prior training and labeled with the reward $\mathcal{R}_{\text{target}}$ from the target MDP. At test time, the agent interacts with a target MDP defined by $\mathcal{M}_{\text{target}} = (\mathcal{S}, \mathcal{A}, \mathcal{P}_{\text{target}}, \mathcal{R}_{\text{target}}, \rho_0, \gamma)$, which presents an aspect of novelty not present in any of the prior MDPs, in the form of new dynamics $\mathcal{P}_{\text{target}}(s_{t+1} | s_t, a_t)$, which may *change over the course of the test-time trial*. We operate in a single-life deployment setting [4] that aims to maximize $J = \sum_{t=0}^h \gamma^t \mathcal{R}(s_t)$, where h is the trial horizon, which may be ∞ . The agent needs to complete the desired task in this target MDP in a single life without any additional supervision or human intervention by effectively selecting and adapting the prior behaviors to the situation at hand.

Off-policy reinforcement learning (RL) algorithms train a parametric Q-function, represented as $Q_\theta(s, a)$, via iterative applications of the Bellman optimality operator, expressed as

$$B^*Q(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim P(s'|s, a)} \left[\max_{a'} Q(s', a') \right].$$

In actor-critic frameworks, a separate policy is trained to maximize Q-values. These algorithms alternate between policy evaluation, which involves the Bellman operator $B^\pi Q = r + \gamma P^\pi Q$, and policy improvement, where the policy $\pi(a|s)$ is updated to maximize expected Q-values. We use a state-of-the-art off-policy actor-critic algorithm RLPD [60] as our base algorithm for pre-training and fine-tuning, which builds on regularized soft actor-critic [61].

IV. ROBUST AUTONOMOUS MODULATION

We now present our method, Robust Autonomous Modulation (ROAM), which fine-tunes value functions with an additional loss and provides a mechanism for choosing among them, so that at deployment time, the agent can quickly react to its current situation *at every timestep* by honing in on the most appropriate behavior from its prior behaviors. Our key observation is that with proper regularization, value functions provide a good indication of how well different behaviors will perform, so we can leverage them to quickly identify appropriate behaviors, which circumvents the need to learn a separate meta-controller or adaptation module.

a) Behavior Modulation using Value Functions.

The core idea of our method is to directly utilize the expressive power of value functions for adaptive behavior selection, as they inherently contain detailed information about potential rewards associated with different states for each behavior. We propose to use value functions of the behaviors

to select the most appropriate behavior for a given state—namely, at each timestep during deployment, choosing one of the behaviors with the high values at that state. Since we already have access to value functions from pre-training the behaviors, this approach does not require any additional training or data collection. Using value functions as the proxy selector also gives much more versatility to the selection mechanism, which can be flexibly controlled on the go by updating the value functions of different behaviors. However, naively using the pre-trained value functions may not lead to high-reward behaviors, due to overestimation of the value functions on new states, as the pre-trained Q-functions may not generalize well to OOD situations. Recent studies in offline RL, for example, have observed an overestimation bias [62] due to OOD training-time actions, which can lead to poor performance when deploying the policy in a new environment. To mitigate these issues, works in offline RL have proposed a number of various modifications aimed at regularizing either the policy or the value function [63]–[65]. Although our setting is different, as we deal with OOD states, we face a similar problem of poor generalization of the value functions of the prior behaviors to unfamiliar situations. In the following section, we describe how we can conservatively regularize the value functions to improve their generalization.

b) Fine-Tuning Value Functions with ROAM.

We desire value functions that accurately reflect the expected reward of using a behavior at a given state. Thus, we would like to fine-tune the value functions of the behaviors to minimize overestimation of the values at states for which a different behavior is more suitable. We can do so by incentivizing the value functions of the behaviors to be higher for states that are visited by that behavior and lower for states visited by other behaviors.

Given a set of prior behaviors \mathcal{B} with policies π_i and critics Q_i , we first fine-tune the value functions of the behaviors with an additional cross-entropy loss on top of the Bellman error that takes in the values at a given state as the logits. Values at a given state s for behavior i are obtained by averaging $Q_i(s, a)$ over $N = 5$ sampled actions $a \sim \pi_i(\cdot | s)$. More formally, with each prior data buffer \mathcal{D}_i , we fine-tune the critic $Q_i(s, a)$ of each behavior i with the following update:

$$\begin{aligned} \mathcal{L}_{\text{fine-tune}} &= (1 - \beta) \mathcal{L}_{\text{Bellman}} + \beta \mathcal{L}_{\text{cross-entropy}} \\ &= (1 - \beta) \sum_i \mathbb{E}_{s, a, s' \sim \mathcal{D}_i} \left[(r(s, a) \right. \\ &\quad \left. + \gamma \mathbb{E}_{a' \sim \pi_i(a'|s')} Q_i(s', a') - Q_i(s, a))^2 \right] \\ &\quad + \beta \sum_j \mathbb{E}_{s \sim \mathcal{D}_j} \left[-\log \frac{\exp(V_j(s))}{\sum_{k=1}^n \exp(V_k(s))} \right], \end{aligned} \quad (1)$$

where $0 < \beta < 1$ is a hyperparameter, $V_i(s) = \mathbb{E}_{a \sim \pi_i(a|s)} [Q_i(s, a)]$ is the average value of behavior i at state s , and \mathcal{D}_i is a replay buffer collected by behavior i . Consider the derivative of the cross-entropy term with respect to the value functions $V_i(s)$ and $V_k(s)$, where $i \neq k$, for a state s in \mathcal{D}_i :

$$\frac{\partial \mathcal{L}_{\text{cross-entropy}}}{\partial V_i(s)} = \frac{\exp(V_i(s))}{\sum_{j=1}^n \exp(V_j(s))} - 1 < 0,$$




Fig. 2: **Robust Autonomous Modulation (ROAM)**. During the initial fine-tuning phase of ROAM, we fine-tune each behavior using its existing data buffer and standard Bellman error, with an additional cross-entropy loss between the softmax values of all behaviors and the behavior index (as the label) from which the state was visited. Then at test-time, at each time-step, we sample from the softmax distribution of the behaviors’ values at the current state and execute the sampled policy.

$$\frac{\partial \mathcal{L}_{\text{cross-entropy}}}{\partial V_k(s)} = \frac{\exp(V_k(s))}{\sum_{j=1}^n \exp(V_j(s))} > 0.$$

So when minimizing the cross-entropy loss, the value function $V_i(s)$ will be pushed up (since its derivative is negative), and $V_k(s)$ for $k \neq i$ will be pushed down. Thus, the cross-entropy loss term in Equation 1 pushes up the value functions of the behaviors for states that are visited by that behavior and pushes down for states that are visited by other behaviors. The value functions are then less likely to overestimate at OOD states, enabling the behaviors to specialize in different parts of the state space, which will help us at test time to better infer an appropriate behavior from the current state.

c) Full Procedure and Single-Life Deployment.

To summarize the full procedure of ROAM, we are given a set of policies π_i and critics Q_i , and a set of prior data buffers \mathcal{D}_i for each behavior, which have been relabeled with the target MDP reward. Alternatively, this can be relaxed and we can assume that we are given a set of prior data buffers \mathcal{D}_i for each behavior, and we can train the policies π_i and critics Q_i using these buffers with offline RL. We then fine-tune the value functions of the behaviors $V_{i,\text{orig}}$ with the additional cross-entropy loss term in Equation 1 to obtain the final value functions V_i .

Then during each timestep at test time, we sample a behavior from the softmax distribution given by the behaviors’ values V_i at the current state. Formally, given the current state s , we sample an action $a_i \sim \pi_i(a|s)$ from behavior i with probability proportional to $\exp(V_i(s))$. The transition (s_t, a_t, r_t, s_{t+1}) is then added to the online buffer $\mathcal{D}_{\text{online}}^i$ for behavior i and the critic V_i and policy π_i and the original value function $V_{i,\text{orig}}$ are fine-tuned using data from $\mathcal{D}_{\text{online}}^i$. We note that the fine-tuned value functions do not affect the policies themselves and only affect which policies are deployed at single-life time. The ROAM fine-tuning objective and single-life deployment are depicted in Figure 2 and the full algorithm is summarized in the Appendix on our project website. In the Appendix, we additionally provide theoretical analysis of ROAM to show that the additional cross-entropy loss in ROAM will lead to more suitable behaviors being chosen.

V. EXPERIMENTAL RESULTS

In this section, we evaluate the performance of ROAM empirically and assess how effectively it can adapt on-the-fly

to new situations. Concretely, we aim to answer the following questions: (1) In simulated and real-world settings, how does ROAM compare to existing methods when given diverse prior behaviors/data and deployed in novel situations? (2) How does the additional cross-entropy term in the loss function of ROAM contribute to more successful utilization of the prior behaviors? In the remainder of this section, we describe our experimental setup and present our results on both a simulated and a real-world Go1 quadrupedal robot. For qualitative video results, see our project webpage: <https://anniesch.github.io/adapt-on-the-go/>.

General Experimental Setup. We use the setting of locomotion deployment to evaluate ROAM, as it is a challenging setting for adaptation, where the agent may naturally face a variety of different situations and must adapt its walking behavior on-the-fly without any additional supervision or human intervention. We use a Go1 quadrupedal robot from Unitree and MuJoCo [66] for simulation. We implemented all methods on top of the same state-of-the-art implementation of SAC from [67] as the base learning approach, with regularization additions following DroQ [68], and RLPD [60] for methods that do online fine-tuning. We include details and hyperparameters on our website.

Comparisons. We evaluate ROAM along with the following prior methods: (1) RLPD Fine-tuning [60], where we fine-tune a single policy using all the data from the prior behaviors with RLPD; (2) RMA [19], which trains a base policy and adaptation module that estimates environment info; (3) High-level Classifier (HLC), which trains a classifier on the data buffers of the pre-trained behaviors and uses it to select which behavior to use at a given state, as a representative method for those that train an additional behavior selection network, similar to work by [69]. We additionally consider an ablation, ROAM-NoFT, which uses the values of the prior behaviors to choose among behaviors but does not fine-tune with the additional cross-entropy loss. We give RMA access to unlimited online episode rollouts in each of the prior MDPs during pre-training, while all other methods use the same set of offline data and prior behaviors that are pre-trained in the prior MDPs.

A. Selecting Relevant Behaviors in Simulation

Setup. In our simulation experiments, we evaluate in two separate settings. The first setting simulates a situation where




Fig. 3: **Results on the simulated Go1 robot.** In both evaluation settings, ROAM is over 2x as efficient as all comparisons in terms of both average and median number of steps taken to complete the task.

different joints become damaged or stuck during the robot’s lifetime. It uses 9 prior behaviors: each is a different limping behavior with a different joint frozen. In the single life, the agent must walk a total distance of 10 meters, and every 100 steps, one of the 3 remaining joints *not covered in the prior data* is frozen, and the agent must adapt its walking behavior to continue walking. The second setting simulates a situation where the robot encounters different friction levels on its different feet due to variation in terrain. It uses 4 different prior behaviors, each of which is trained with one of the 4 feet having low friction. During the single life, every 50 steps, the friction of one or two of the feet is changed to be lower than in the prior behaviors. To collect the prior behaviors, we train each behavior for 250k steps (first setting) or 50k steps (second setting) in the corresponding MDP, and use the last 40k steps as the prior data. The agent is given a maximum of 10k steps to complete the task, and if it does not complete the task within this time, it is considered an unsuccessful trial.

Results. As seen in Figure 3, ROAM outperforms all other methods in all three metrics of average and median number of steps taken to complete the task. In particular, in both settings, ROAM completes the task *more than 2 times faster*, in terms of average number of timesteps, compared to the next best method. Both RLPD fine-tuning and RMA struggle on both evaluation settings, especially the stiffness setting, demonstrating the importance of adapting in the space of behaviors rather than the space of actions for more efficient adaptation. RLPD and RMA perform better in the friction evaluation, where a single policy can still somewhat adapt to the various situations in the single life. On the other hand, HLC and ROAM-NoFT both struggle in the friction eval suite, demonstrating the importance of the additional cross-entropy term in the loss function, encouraging greater behavior specialization in different regions of the state space. These two methods perform better when the behaviors are already more distinguished, as in the limping eval suite, but they still struggle to adapt as efficiently as ROAM.

In Figures 4 and 5, we provide some additional empirical analysis of ROAM. First, in the stiffness evaluation, we plot the percent of steps where different methods select a relevant behavior during test-time deployment, where a held-out joint is frozen and a relevant behavior is one where an adjacent joint on the same leg is frozen or the same joint on an adjacent leg is frozen, and we see that ROAM is able to choose the most

relevant behavior on average significantly more frequently than HLC and ROAM-noFT, which often select a behavior that is not relevant to the current situation. In addition, we record the average values of the different behaviors in states visited by that behavior (ID) vs states visited by other behaviors (OOD), before and after fine-tuning with the additional cross-entropy term for ROAM. We find that ROAM is able to effectively maintain high values of the behavior in familiar states, while decreasing the value of the behavior in unfamiliar OOD states.

B. Adapting on-the-Go1 in the Real World

Setup. On the Go1 quadruped robot, we evaluate ROAM in a setting where we have a fixed set of five prior behaviors: walking, and four different behaviors where each of the legs has a joint frozen. We pre-train a base walking behavior in 18k steps and train the other behaviors by fine-tuning the walking behavior for an additional 3k steps with one of the joints frozen, all from scratch in the real world using the system from [67]. During single-life deployment, we evaluate on the following three tasks: (1) Heavy Luggage: the robot must walk from a starting line to a finish line, while pulling a box that is 6.2 kg (13.6 lb) attached to one of the back legs. (2) Dynamic Luggage Load: the robot must walk from a starting line to a finish line, while adapting on-the-fly to a varying amount of weight between 2.36 kg (5.2 lb) and 4.2 kg (9.2 lb). We standardize each trial by adding and removing weight at the same distance from the start position. (3) Roller Skates: we fit the robot’s front two feet into roller skates, and the robot must adapt to its behavior to slide its forward legs and push off its back legs in order to walk to the end line. We report the average wall clock time in seconds and number of falls or readjustments needed to complete the task in a single life across 3 trials for each method. The tasks are shown in Figure 6, along with the locations of the robot before its first fall or readjustment for each method for each trial.

Results. Although none of the prior behaviors are trained to handle these specific test-time scenarios, the robot can leverage parts of the prior behaviors to complete the task. As shown in Table I, ROAM significantly outperforms using a high-level classifier (HLC) as well as the baseline walking policy in terms of both average wall clock time and number of falls or readjustments at single-life time on all three real-world tasks. Qualitatively, the other methods have trouble pulling luggage consistently forward, whereas our method often chooses the




Fig. 4: **Single-life Behavior Selection Accuracy.** The percent of steps where different methods select a relevant behavior for the current situation. ROAM is able to choose a relevant behavior significantly more often on average when adapting to test-time situations than HLC and ROAM-noFT.

		Before FT	After FT
Changing Stiffness	Avg ID Values	779.2	778.7
	Avg OOD Values	746.4	542.9
Changing Frictions	Avg ID Values	1243.4	1247.6
	Avg OOD Values	1217.6	1192.8

Fig. 5: **Effect of the Cross-Entropy Loss on ID and OOD values.** The average values of the different behaviors in states visited by that behavior (ID) vs states visited by other behaviors (OOD), before and after fine-tuning with the additional cross-entropy term for ROAM. ROAM is able to maintain high values of the behavior in ID states, while decreasing the value of the behavior in OOD states.

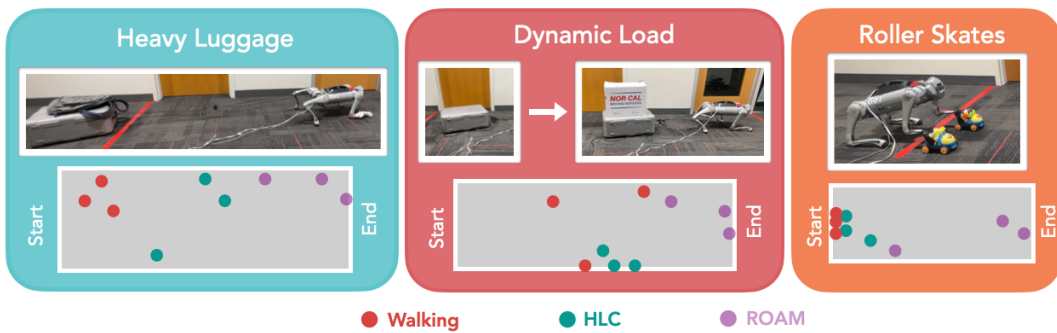


Fig. 6: **Real-world single-life tasks.** We evaluate on: (1) pulling a load of heavy luggage 6.2 kg (13.6 lb), (2) pulling luggage where the weight dynamically changes between 2.36 kg (5.2 lb) and 4.2 kg (9.2 lb), and (3) moving forward with roller skates on the robot’s front two feet. For each trial, for each method, we also show the locations of the robot before its first fall or readjustment (Red is Walking, Blue is HLC, and Purple is ROAM).

	Heavy Luggage		Dynamic Luggage Load		Roller Skates	
	Avg. Time (s) ↓	Falls ↓	Avg. Time (s) ↓	Falls ↓	Avg. Time (s) ↓	Falls ↓
Walking	45.3	2.3	32	1	NC	NC
HLC	42.7	3	28.3	1.3	62.3	2.7
ROAM (ours)	25.7	0.7	24.3	0.3	27.3	1

TABLE I: **Results on the real Go1 robot on 3 different tasks:** On all 3 tasks, across 3 trials for each method, ROAM significantly outperforms both comparisons in terms of both average wall clock time (s) and number of falls or readjustments needed to complete the task in a single life. NC (no complete) indicates that the task was not able to be successfully completed at all with the given method.

behavior where a joint is frozen on the leg with the luggage attached, as this behavior uses the robot’s other three legs to pull itself forward more effectively. The other methods struggle particularly on the roller skates task, which has drastically different dynamics from typical walking and especially relies on choosing relevant behaviors that heavily use the back legs. As seen in Figure 6, for all three tasks, HLC and the standard walking policy often fall or need to be readjusted very early in each single-life trial, whereas ROAM gets much closer to the finish line and often even completes the task without any falls or readjustments.

VI. CONCLUSION AND FUTURE WORK

We introduced Robust Autonomous Modulation (ROAM), which enables agents to rapidly adapt to changing, out-of-distribution circumstances during deployment. Our contribution lies in offering a principled, efficient way for agents to leverage pre-trained behaviors when adapting on-the-fly. On

simulated and complex real-world tasks with a Go1 quadruped robot, our method achieves over 2x efficiency in adapting to new situations compared to existing methods. While ROAM offers significant advances in enabling agents to adapt to out-of-distribution scenarios, one current limitation lies in the dependency on the range of pre-trained behaviors; some scenarios may simply be too far out-of-distribution compared to the available prior behaviors. Future work could explore integrating ROAM into a lifelong learning framework, allowing agents to continuously expand their repertoire of behaviors and increasing their adaptability to more unforeseen situations.

ACKNOWLEDGMENTS

We would like to thank the IRIS and RAIL labs for helpful feedback on the project. This research was supported by an NSF Graduate Research Fellowship, Apple, and ONR grants N00014-20-1-2675 and N00014-21-1-2685.

REFERENCES

- [1] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, no. 7553, pp. 503–507, 2015.
- [2] X. Song, Y. Yang, K. Choromanski, K. Caluwaerts, W. Gao, C. Finn, and J. Tan, "Rapidly adaptable legged robots via evolutionary meta-learning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 3769–3776.
- [3] R. Julian, B. Swanson, G. Sukhatme, S. Levine, C. Finn, and K. Hausman, "Never stop learning: The effectiveness of fine-tuning in robotic reinforcement learning," 2020. [Online]. Available: <https://arxiv.org/abs/2004.10190>
- [4] A. Chen, A. Sharma, S. Levine, and C. Finn, "You only live once: Single-life reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 14784–14797, 2022.
- [5] M. Cutler, T. J. Walsh, and J. P. How, "Reinforcement learning with multi-fidelity simulators," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 3888–3895.
- [6] A. Rajeswaran, S. Ghotra, B. Ravindran, and S. Levine, "Epopt: Learning robust neural network policies using model ensembles," *arXiv preprint arXiv:1610.01283*, 2016.
- [7] F. Sadeghi and S. Levine, "Cad2rl: Real single-image flight without a single real image," *arXiv preprint arXiv:1611.04201*, 2016.
- [8] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 23–30.
- [9] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 3803–3810.
- [10] J. Tan, T. Zhang, E. Coumans, A. Iscen, Y. Bai, D. Hafner, S. Bohez, and V. Vanhoucke, "Sim-to-real: Learning agile locomotion for quadruped robots," *arXiv preprint arXiv:1804.10332*, 2018.
- [11] W. Yu, V. C. Kumar, G. Turk, and C. K. Liu, "Sim-to-real transfer for biped locomotion," in *2019 IEEE/RSJ international conference on intelligent robots and systems (iros)*. IEEE, 2019, pp. 3503–3510.
- [12] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas *et al.*, "Solving rubik's cube with a robot hand," *arXiv preprint arXiv:1910.07113*, 2019.
- [13] Z. Xie, X. Da, M. Van de Panne, B. Babich, and A. Garg, "Dynamics randomization revisited: A case study for quadrupedal locomotion," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 4955–4961.
- [14] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal, "Rapid locomotion via reinforcement learning," *arXiv preprint arXiv:2205.02824*, 2022.
- [15] T. Haarnoja, B. Moran, G. Lever, S. H. Huang, D. Tirumala, M. Wulfmeier, J. Humpalik, S. Tunyasuvunakool, N. Y. Siegel, R. Hafner *et al.*, "Learning agile soccer skills for a bipedal robot with deep reinforcement learning," *arXiv preprint arXiv:2304.13653*, 2023.
- [16] W. Yu, J. Tan, C. K. Liu, and G. Turk, "Preparing for the unknown: Learning a universal policy with online system identification," *RSS*, 2017.
- [17] G. Ji, J. Mun, H. Kim, and J. Hwangbo, "Concurrent training of a control policy and a state estimator for dynamic and robust legged locomotion," *IEEE Robotics and Automation Letters*, 2022.
- [18] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, p. eabc5986, 2020.
- [19] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "Rma: Rapid motor adaptation for legged robots," *arXiv preprint arXiv:2107.04034*, 2021.
- [20] X. B. Peng, E. Coumans, T. Zhang, T.-W. Lee, J. Tan, and S. Levine, "Learning agile robotic locomotion skills by imitating animals," *arXiv preprint arXiv:2004.00784*, 2020.
- [21] W. Yu, J. Tan, Y. Bai, E. Coumans, and S. Ha, "Learning fast adaptation with meta strategy optimization," *IEEE Robotics and Automation Letters*, 2020.
- [22] Z. Fu, X. Cheng, and D. Pathak, "Deep whole-body control: learning a unified policy for manipulation and locomotion," in *Conference on Robot Learning*, 2022.
- [23] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, 2022.
- [24] A. Agarwal, A. Kumar, J. Malik, and D. Pathak, "Legged locomotion in challenging terrains using egocentric vision," in *Conference on Robot Learning*, 2022.
- [25] Z. Zhuang, Z. Fu, J. Wang, C. G. Atkeson, S. Schertfeger, C. Finn, and H. Zhao, "Robot parkour learning," in *7th Annual Conference on Robot Learning*, 2023.
- [26] R. Yang, G. Yang, and X. Wang, "Neural volumetric memory for visual locomotion control," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [27] J. X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Z. Leibo, R. Munos, C. Blundell, D. Kumaran, and M. Botvinick, "Learning to reinforcement learn," *arXiv preprint arXiv:1611.05763*, 2016.
- [28] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel, "RL²: Fast reinforcement learning via slow reinforcement learning," *arXiv preprint arXiv:1611.02779*, 2016.
- [29] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 1126–1135.
- [30] A. Nagabandi, I. Clavera, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn, "Learning to adapt in dynamic, real-world environments through meta-reinforcement learning," *arXiv preprint arXiv:1803.11347*, 2018.
- [31] R. Houthoofd, Y. Chen, P. Isola, B. Stadie, F. Wolski, O. Jonathan Ho, and P. Abbeel, "Evolved policy gradients," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [32] J. Rothfuss, D. Lee, I. Clavera, T. Asfour, and P. Abbeel, "Promp: Proximal meta-policy search," *arXiv preprint arXiv:1810.06784*, 2018.
- [33] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell, "Meta-learning with latent embedding optimization," *arXiv preprint arXiv:1807.05960*, 2018.
- [34] R. Mendonca, X. Geng, C. Finn, and S. Levine, "Meta-reinforcement learning robust to distributional shift via model identification and experience relabeling," *arXiv preprint arXiv:2006.07178*, 2020.
- [35] N. Hansen, R. Jangir, Y. Sun, G. Alenyà, P. Abbeel, A. A. Efros, L. Pinto, and X. Wang, "Self-supervised policy adaptation during deployment," *arXiv preprint arXiv:2007.04309*, 2020.
- [36] T. Yoneda, G. Yang, M. R. Walter, and B. Stadie, "Invariance through inference," *arXiv preprint arXiv:2112.08526*, 2021.
- [37] K. Khetarpal, M. Riemer, I. Rish, and D. Precup, "Towards continual reinforcement learning: A review and perspectives," *arXiv preprint arXiv:2012.13490*, 2020.
- [38] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell, "Progressive neural networks," *arXiv preprint arXiv:1606.04671*, 2016.
- [39] B. Eysenbach, S. Asawa, S. Chaudhari, S. Levine, and R. Salakhutdinov, "Off-dynamics reinforcement learning: Training for transfer with domain classifiers," *arXiv preprint arXiv:2006.13916*, 2020.
- [40] A. Xie, J. Harrison, and C. Finn, "Deep reinforcement learning amidst lifelong non-stationarity," *arXiv preprint arXiv:2006.10701*, 2020.
- [41] A. Xie and C. Finn, "Lifelong robotic reinforcement learning by retaining experiences," *arXiv preprint arXiv:2109.09180*, 2021.
- [42] K. Gregor, D. J. Rezende, and D. Wierstra, "Variational intrinsic control," *arXiv preprint arXiv:1611.07507*, 2016.
- [43] J. Achiam, H. Edwards, D. Amodei, and P. Abbeel, "Variational option discovery algorithms," *arXiv preprint arXiv:1807.10299*, 2018.
- [44] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine, "Diversity is all you need: Learning skills without a reward function," *arXiv preprint arXiv:1802.06070*, 2018.
- [45] O. Nachum, S. Gu, H. Lee, and S. Levine, "Near-optimal representation learning for hierarchical reinforcement learning," *arXiv preprint arXiv:1810.01257*, 2018.
- [46] A. Sharma, S. Gu, S. Levine, V. Kumar, and K. Hausman, "Dynamics-aware unsupervised discovery of skills," *arXiv preprint arXiv:1907.01657*, 2019.
- [47] K. Baumli, D. Warde-Farley, S. Hansen, and V. Mnih, "Relative variational intrinsic control," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 35, no. 8, 2021, pp. 6732–6740.
- [48] M. Laskin, H. Liu, X. B. Peng, D. Yarats, A. Rajeswaran, and P. Abbeel, "Cic: Contrastive intrinsic control for unsupervised skill discovery," *arXiv preprint arXiv:2202.00161*, 2022.
- [49] S. Park and S. Levine, "Predictable mdp abstraction for unsupervised model-based rl," *arXiv preprint arXiv:2302.03921*, 2023.
- [50] P.-L. Bacon, J. Harb, and D. Precup, "The option-critic architecture," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 31, no. 1, 2017.
- [51] X. B. Peng, M. Chang, G. Zhang, P. Abbeel, and S. Levine, "Mcp: Learning composable hierarchical control with multiplicative compositional policies," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

- [52] Y. Lee, J. Yang, and J. J. Lim, "Learning to coordinate manipulation skills via skill behavior diversification," in *International conference on learning representations*, 2019.
- [53] M. Sharma, J. Liang, J. Zhao, A. LaGrassa, and O. Kroemer, "Learning to compose hierarchical object-centric controllers for robotic manipulation," *arXiv preprint arXiv:2011.04627*, 2020.
- [54] R. Strudel, A. Pashevich, I. Kalevtykh, I. Laptev, J. Sivic, and C. Schmid, "Learning to combine primitive skills: A step towards versatile robotic manipulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 4637–4643.
- [55] O. Nachum, S. S. Gu, H. Lee, and S. Levine, "Data-efficient hierarchical reinforcement learning," *Advances in neural information processing systems*, vol. 31, 2018.
- [56] R. Chitnis, S. Tulsiani, S. Gupta, and A. Gupta, "Efficient bimanual manipulation using learned task schemas," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1149–1155.
- [57] K. Pertsch, Y. Lee, Y. Wu, and J. J. Lim, "Guided reinforcement learning with learned skills," *arXiv preprint arXiv:2107.10253*, 2021.
- [58] M. Dalal, D. Pathak, and R. R. Salakhutdinov, "Accelerating robotic reinforcement learning via parameterized action primitives," *Advances in Neural Information Processing Systems*, vol. 34, pp. 21 847–21 859, 2021.
- [59] S. Nasiriany, H. Liu, and Y. Zhu, "Augmenting reinforcement learning with behavior primitives for diverse manipulation tasks," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 7477–7484.
- [60] P. J. Ball, L. Smith, I. Kostrikov, and S. Levine, "Efficient online reinforcement learning with offline data," *arXiv preprint arXiv:2302.02948*, 2023.
- [61] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International conference on machine learning*. PMLR, 2018, pp. 1861–1870.
- [62] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline reinforcement learning: Tutorial, review, and perspectives on open problems," *arXiv preprint arXiv:2005.01643*, 2020.
- [63] A. Kumar, A. Zhou, G. Tucker, and S. Levine, "Conservative q-learning for offline reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1179–1191, 2020.
- [64] T. Yu, G. Thomas, L. Yu, S. Ermon, J. Y. Zou, S. Levine, C. Finn, and T. Ma, "Mopo: Model-based offline policy optimization," *Advances in Neural Information Processing Systems*, vol. 33, pp. 14 129–14 142, 2020.
- [65] T. Yu, A. Kumar, R. Rafailov, A. Rajeswaran, S. Levine, and C. Finn, "Combo: Conservative offline model-based policy optimization," *Advances in neural information processing systems*, vol. 34, pp. 28 954–28 967, 2021.
- [66] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 5026–5033.
- [67] L. Smith, I. Kostrikov, and S. Levine, "A walk in the park: Learning to walk in 20 minutes with model-free reinforcement learning," *arXiv preprint arXiv:2208.07860*, 2022.
- [68] T. Hiraoka, T. Imagawa, T. Hashimoto, T. Onishi, and Y. Tsuruoka, "Dropout q-functions for doubly efficient reinforcement learning," *arXiv preprint arXiv:2110.02034*, 2021.
- [69] L. Han, Q. Zhu, J. Sheng, C. Zhang, T. Li, Y. Zhang, H. Zhang, Y. Liu, C. Zhou, R. Zhao *et al.*, "Lifelike agility and play on quadrupedal robots using reinforcement learning and generative pre-trained models," *arXiv preprint arXiv:2308.15143*, 2023.

A. Theoretical Analysis

In this section, we theoretically analyze ROAM to show that the additional cross-entropy loss in ROAM will lead to more suitable behaviors being chosen at each timestep. In particular, ROAM rescales the value functions of the behaviors so that they are less likely to overestimate in states that are out of distribution for that behavior. Our main result, given in Theorem .1, is that with ROAM, for some weight $\beta > 0$ on the cross-entropy term, at a given state, ROAM constrains each behavior's value to be lower than the value of behaviors for which that state appears more frequently. This theorem gives us *conservative generalization* by reducing value overestimation in unfamiliar states—specifically, if at least one behavior is familiar with the current state, our chosen behavior will not have worse performance than the most familiar behavior.

Theorem .1. *Let $p_i(s)$ denote the state visitation probability for a behavior i at state s . For any state s that is out of distribution for behavior i and is in distribution for another behavior j , i.e. $p_i(s) \ll p_j(s)$, if $0 < \beta < 1$ is chosen to be large enough, then the value of behavior i learned by ROAM will be bounded above compared to value of behavior j , i.e., $V_i(s) \leq V_j(s)$.*

This theoretical result holds for all local extrema. Under perfect optimization, ROAM will reach a local minimum at convergence and the result would hold in this case. Our analysis implicitly assumes a simplified tabular setting in the absence of approximation and sampling errors, as done in prior theoretical analysis of deep RL algorithms, e.g. some of the analysis in [63]. We take the full gradient of the TD error in the proof although in practice, only the semi-gradient is used during optimization. Regardless, the optimal value function is still a minimum of $\mathcal{L}_{\text{fine-tune}}$ and should satisfy the Bellman equation in the loss function at convergence. We also assume that for all states s , we have $V_{\text{freq}}(s) - V_i(s) = (R_{\text{freq}}(s) + \gamma \mathbb{E}_{s'} V_{\text{freq}}(s')) - (R_i(s) + \gamma \mathbb{E}_{s'} V_i(s')) < \eta < \infty$.

Theorem .2. *Let $p_i(s)$ denote the state visitation probability for a behavior b_i at state s . For any state s that is out of distribution for behavior b_i and is in distribution for another behavior b_j , i.e. $p_i(s) \ll p_j(s)$, if $1 > \beta > 0$ is chosen to be large enough, then the value of behavior i learned by ROAM will be bounded above compared to value of behavior b_j , i.e. $V_i(s) \leq V_j(s)$.*

Proof. We optimize the following loss function:

$$\mathcal{L}_{\text{fine-tune}} = (1 - \beta) \sum_i \sum_{s \sim \mathcal{D}_i} p_i(s) \left[(R_i(s) + \gamma \mathbb{E}_{s'} V_i(s') - V_i(s))^2 \right] + \beta \sum_j \sum_{s \sim \mathcal{D}_j} p_j(s) \left[-V_j(s) + \log \sum_{k=1}^n \exp(V_k(s)) \right].$$

Taking the derivative with respect to $V_i(s)$, we have:

$$\begin{aligned} \frac{\partial \mathcal{L}_{\text{fine-tune}}}{\partial V_i(s)} &= 2(1 - \beta) p_i(s) (\gamma p(s|s, a) - 1) (R_i(s) + \gamma \mathbb{E}_{s'} V_i(s') - V_i(s)) \\ &+ \beta \left[p_i(s) \left(-1 + \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \right) + \sum_{j \neq i} p_j(s) \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] \\ &= 2(1 - \beta) p_i(s) (\gamma P_i(s|s, a) - 1) (R_i(s) + \gamma \mathbb{E}_{s'} V_i(s') - V_i(s)) + \beta \left[-p_i(s) + \sum_j p_j(s) \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \right]. \end{aligned}$$

Setting to 0, we have the following characterization of any convergence point:

$$\begin{aligned} V_i(s) &= (R_i(s) + \gamma \mathbb{E}_{s'} V_i(s')) + \frac{\beta}{2(1 - \beta) p_i(s) (\gamma P_i(s|s, a) - 1)} \left[-p_i(s) + \sum_j p_j(s) \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] \\ &= (R_i(s) + \gamma \mathbb{E}_{s'} V_i(s')) + \frac{\beta}{2(1 - \beta) (1 - \gamma P_i(s|s, a))} \left[1 - \sum_j \frac{p_j(s)}{p_i(s)} \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \right]. \end{aligned}$$

Consider a state s where $p_{\text{freq}}(s) \gg p_i(s)$ for some behavior $b_{\text{freq}} \neq b_i$. We want to show that in such a state $V_{\text{freq}}(s) > V_i(s)$. Then $\frac{p_{\text{freq}}(s)}{p_i(s)}$ will dominate in the last term, so

$$\sum_j \frac{p_j(s)}{p_i(s)} \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \gg 1. \quad (2)$$

Then comparing the convergence point values of $V_i(s)$ and $V_{\text{freq}}(s)$, we have

$$\begin{aligned} V_{\text{freq}}(s) - V_i(s) &= (R_{\text{freq}}(s) + \gamma \mathbb{E}_{s'} V_{\text{freq}}(s')) + \frac{\beta}{2(1-\beta)(1-\gamma P_{\text{freq}}(s|s,a))} \left[1 - \sum_j \frac{p_j(s)}{p_{\text{freq}}(s)} \frac{\exp(V_{\text{freq}}(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] \\ &\quad - \left[(R_i(s) + \gamma \mathbb{E}_{s'} V_i(s')) + \frac{\beta}{2(1-\beta)(1-\gamma P_i(s|s,a))} \left[1 - \sum_j \frac{p_j(s)}{p_i(s)} \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] \right] \\ &= ((R_{\text{freq}}(s) + \gamma \mathbb{E}_{s'} V_{\text{freq}}(s')) - (R_i(s) + \gamma \mathbb{E}_{s'} V_i(s'))) \\ &\quad + \frac{\beta}{1-\beta} \left(C_{\text{freq}} \left[1 - \sum_j \frac{p_j(s)}{p_{\text{freq}}(s)} \frac{\exp(V_{\text{freq}}(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] - C_i \left[1 - \sum_j \frac{p_j(s)}{p_i(s)} \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] \right), \end{aligned}$$

where $C_{\text{freq}} = \frac{1}{2(1-\gamma P_{\text{freq}}(s|s,a))} > 0$ and $C_i = \frac{1}{2(1-\gamma P_i(s|s,a))} > 0$. Thus, using Equation 2, because

$$1 - \sum_j \frac{p_j(s)}{p_i(s)} \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \ll 0,$$

the term

$$\left(C_{\text{freq}} \left[1 - \sum_j \frac{p_j(s)}{p_{\text{freq}}(s)} \frac{\exp(V_{\text{freq}}(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] - C_i \left[1 - \sum_j \frac{p_j(s)}{p_i(s)} \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] \right) > 0.$$

Hence, for some $0 < \beta < 1$, we have

$$\begin{aligned} V_{\text{freq}}(s) - V_i(s) &= ((R_{\text{freq}}(s) + \gamma \mathbb{E}_{s'} V_{\text{freq}}(s')) - (R_i(s) + \gamma \mathbb{E}_{s'} V_i(s'))) \\ &\quad + \frac{\beta}{1-\beta} \left(C_{\text{freq}} \left[1 - \sum_j \frac{p_j(s)}{p_{\text{freq}}(s)} \frac{\exp(V_{\text{freq}}(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] - C_i \left[1 - \sum_j \frac{p_j(s)}{p_i(s)} \frac{\exp(V_i(s))}{\sum_{k=1}^n \exp(V_k(s))} \right] \right) > 0. \end{aligned}$$

Thus, for some for some $0 < \beta < 1$, we have $V_i(s) < V_{\text{freq}}(s)$. □

As a result, for any states s that are out of distribution for behavior i and in distribution for a different behavior j , if we choose β large enough, the value learned will not overestimate the value compared to behavior j . Thus, at each time step, if one or more behaviors are familiar with the current state, the performance of the chosen behavior will not be much worse than its value function estimate. In this manner, ROAM adjusts value estimates based on degree of familiarity, mitigating overestimation risks. The ability to adjust the β parameter offers a flexible framework to optimize for the behavior with the highest value at a given state, which will be at least as suitable as the most familiar behavior.

To illustrate the effect of the cross-entropy loss, consider the following example of choosing between two behaviors i and j at state s , where the true values $V_i^{\text{true}}(s) < V_j^{\text{true}}(s)$. The optimal choice is to choose behavior j and for sake of this example let us choose behavior j if $V_i^{\text{ROAM}}(s) < V_j^{\text{ROAM}}(s)$. There are the following four cases: (1) $p_i(s) < p_j(s)$ and the initial estimated $V_i(s) < V_j(s)$. Then with any $\beta > 0$, the final $V_i^{\text{ROAM}}(s) < V_j^{\text{ROAM}}(s)$; (2) $p_i(s) < p_j(s)$ and the initial estimated $V_i(s) > V_j(s)$. Then by Theorem .2, with large enough $\beta > 0$, the final $V_i^{\text{ROAM}}(s) < V_j^{\text{ROAM}}(s)$; (3) $p_i(s) > p_j(s)$ and the initial estimated $V_i(s) < V_j(s)$. Then as long as β is chosen to be not too large, the final $V_i^{\text{ROAM}}(s) < V_j^{\text{ROAM}}(s)$. (4) $p_i(s) > p_j(s)$ and the initial estimated $V_i(s) > V_j(s)$. This is the only case where ROAM may not be adjusted to work well, but this case poses a difficult situation for any behavior selection method.

B. Algorithm Summary

We summarize ROAM in Algorithms 1 and 2.

C. Additional Empirical Analysis

The cross-entropy term is a regularizer that creates a preference for skills that visit a given state more frequently. However, this is not the only criterion for selecting a skill; it is a regularizer. A skill with higher value is still preferred if its visitation frequency is not too low, and ROAM does not exclusively always just select the most high-frequency behavior. We show this with the following experiment with results in Table II. In the simulated stiffness suite, we held out most of the data from one of the buffers corresponding to one of the behaviors, leaving only 5k transitions compared to the original 40k, and evaluated the agent at test time in an environment suited for that behavior. We find that even with only 5k transitions (compared to 40k for all other behaviors), ROAM selects this less-frequent but suitable behavior the majority of the time, leading to similar overall performance.

Algorithm 1 ROAM FINE-TUNING

```

1: Require:  $\mathcal{D}_i$ , pre-trained critics  $Q_{i,\text{orig}}$ 
2: while not converged do
3:   for all  $i$  in  $1, \dots, N_{\text{behaviors}}$  do
4:     Sample  $(s, a, s', r_{\text{target}}) \sim D_i$ 
5:     Update  $Q_i$  according to Eq. 1
6:   end for
7: end while  $Q_1, \dots, Q_{N_{\text{behaviors}}}$ 

```

Algorithm 2 ROAM SINGLE-LIFE DEPLOYMENT

```

1: Require: Test MDP  $\mathcal{M}_{\text{test}}$ ,  $\mathcal{D}_i$ , policies  $\pi_i$  and fine-tuned critics  $Q_i$ ;
2: Initialize: online replay buffers  $\mathcal{D}_{\text{online}}^i$ ; timestep  $t = 0$ 
3: while task not complete do
4:   Compute values of each behavior  $\{V_i(s_t)\}_1^{N_{\text{behaviors}}}$ 
5:   Sample behavior  $b^*$  according to the distribution
     softmax( $\exp(V_i(s_t))$ ).
6:   Take action  $a_t \sim \pi_{b^*}(a_t|s_t)$ .
7:    $\mathcal{D}_{\text{online}}^{b^*} \leftarrow \mathcal{D}_{\text{online}}^{b^*} \cup \{(s_t, a_t, r_t, s_{t+1})\}$ 
8:    $Q_{b^*,\text{orig}}(s, a), \pi_{b^*} \leftarrow \text{RL}(Q_{b^*,\text{orig}}(s, a), \pi_{b^*}, \mathcal{D}_{\text{online}}^{b^*})$ 
9:   Increment  $t$ 
10: end while

```

# Transitions	% Timesteps Chosen	Avg # Steps
5k	53.2	591.3
40k	78.4	573.8

TABLE II: blackROAM selects high-value behaviors even with lower visitation frequency. We find that even with a much smaller buffer, and therefore lower visitation frequency for many states, ROAM still chooses that behavior when given situations suitable for it.

black We next investigate the sensitivity of the β hyperparameter. We ran ROAM with 4 different values (0.01, 0.1, 0.5, 0.9) of β in each simulated suite and show the performance in Table III. For both evaluations, 3 out of 4 of these values (all except 0.01) outperform all the other baselines.

Additionally, one benefit of ROAM is that the ability to switch between these policies at any timestep allows the agent to adapt to new and unforeseen situations, including those for which no single behavior is optimally suited. However, one hypothetical concern may be that frequent switching of behaviors may lead to suboptimal performance. In Table III, we measure how often behaviors were switched and tried to see if frequency of behavior switches correlates with failure. We found no such correlation. Below, we show the percent of timesteps where the agent decides to switch behaviors, and more frequent switching does not correlate to a higher average number of steps needed to complete the task.

β	Dynamic Friction		Dynamic Stiffness	
	Avg # Steps (\downarrow)	Frequency of Switching	Avg # Steps (\downarrow)	Frequency of Switching
0.01	7610 +- 854	17.20%	2698 +- 844	2.92%
0.1	2082 +- 382	15.63%	1331 +- 263	8.25%
0.5	772 +- 179	11.85%	628 +- 19	12.35%
0.9	1466 +- 534	9.36%	735 +- 54	13.36%

TABLE III: blackSensitivity of β and frequency of behavior switching. We find that a range of β values give strong performance for ROAM.

At test time, ROAM provides the option to use the collected online samples to further fine-tune the policies. We run an additional ablation on this aspect of fine-tuning in our simulated tasks. Below in Table IV, we find that only performing fine-tuning (the RLPD comparison) leads to much slower adaptation. We also add a comparison to ROAM without low-level behavior fine-tuning (ROAM-selection only) in the table below, showing that fine-tuning the low-level behaviors helps but the lion’s share of the efficiency benefit comes from the behavior selection. Thus, the behavior selection mechanism is indeed the main driver of the fast adaptation, but our framework also supports fine-tuning the selected behaviors.

	Dynamic Friction	Dynamic Stiffness
	Avg # Steps (\downarrow)	Avg # Steps (\downarrow)
RLPD	8089 \pm 868	3797 \pm 1029
ROAM (selection only)	1268 \pm 334	1565 \pm 339
ROAM	1018 \pm 225	1199 \pm 439

TABLE IV: Ablation on Behavior-Level Fine-Tuning. ROAM supports fine-tuning the selected behaviors with the additional collected data. The majority of efficiency benefits come from the behavior selection mechanism, though fine-tuning selected behaviors also improves performance.

Finally, in our main experiments in Section V, we range from choosing among 4-9 different prior behaviors. We run an additional experiment combining all the behaviors from both simulated suites, giving 13 total, and in Table V, we find that for each suite, adding additional irrelevant behaviors did not significantly hurt performance and can even improve performance (as is the case for the stiffness eval). We hope future work may explore this at a larger scale.

	Dynamic Friction	Dynamic Stiffness
	Avg # Steps (\downarrow)	Avg # Steps (\downarrow)
ROAM (original, selected skills)	1018 \pm 225	1199 \pm 439
ROAM (all skills)	1141 \pm 115	871 \pm 137

TABLE V: **ROAM with irrelevant behaviors.** Adding irrelevant behaviors to the set of prior behaviors did not significantly hurt performance and sometimes improves it, as seen in the stiffness evaluation.

D. General Experiment Setup Details

As common practice in learning-based quadrupedal locomotion works, we define actions to be PD position targets for the 12 joints, and we use a control frequency of 20 Hz. Actions are centered around the nominal pose, i.e. 0 is standing. We describe the observations for the simulated and real-world experiments below.

We first detail the reward function we use to define the quadrupedal walking task. First, we have a velocity-tracking term defined as follows:

$$r_v(s, a) = 1 - \left| \frac{v_x - v_t}{v_t} \right|^{1.6}$$

where v_t is the target velocity and v_x is the robot’s local, forward linear velocity projected onto the ground plane, i.e., $v_x = v_{\text{local}} \cdot \cos(\phi)$ where ϕ is the root body’s pitch angle. We then have a term $r_{\text{ori}}(s, a)$ that encourages the robot to stay upright. Specifically, we calculate the cosine distance between the 3d vector perpendicular to the robot’s body and the gravity vector $([0, 0, 1])$. We then normalize the term to be between 0 and 1 via:

$$r_{\text{ori}}(s, a) = (0.5 \cdot \text{dist} + 0.5)^2$$

where dist is the cosine distance as described above. We multiply r_v and r_{ori} so as to give reward for tracking velocity proportionally to how well the robot is staying upright. We then have a regularization term r_{qpos} to favor solutions that are close to the robot’s nominal standing pose. This regularization term is calculated as a product of a normalized term per-joint. Below, \hat{q}^j represents the local rotation of joint j of the nominal pose, and q^j represents the robot’s joint,

$$r^{\text{qpos}} = 1 - \prod_j \text{q distance}(\hat{q}^j, q^j)$$

where q distance is between 0 and 1 and decays quadratically until a threshold which is the robot’s action limits. Specifically, we follow the reward structure put forth by tunyasuvunakool2020. These terms comprise the overwhelming majority of weight in the final reward. We also include terms for avoiding undesirable behaviors like rocking or swaying that penalize any angular velocity in the root body’s roll, pitch, and yaw. We also slightly penalize energy consumption and torque smoothness. To encourage a walking gait in particular, we added another regularization term to encourage diagonal shoulder and hip joints to be the same at any given time.

E. Implementation Details and Hyperparameters

We implemented all methods, including ROAM, RMA, and HLC, on top of the same state-of-the art implementation of SAC from smith2022walk as the base learning approach. For all comparisons, we additionally use a high UTD ratio, dropout, and layernorm, following DroQ hiraoka2021dropout, and for methods that do online fine-tuning, we use 50/50 sampling following RLPD ball2023rlpd. We use default hyperparameter values: a learning rate of 3×10^{-4} , an online batch size of 128, and a discount factor of 0.99. The policy and critic networks are MLPs with 2 hidden layers of 256 units each and ReLU activations. For ROAM, we tuned β with values 0.01, 0.1, 0.5, 0.9.

a) Simulated Experiments.

For the simulated experiments, the state space consists of joint positions, joint velocities, torques, IMU (roll, pitch, change in roll, change in pitch), and normalized foot forces for a total of 44 dimensions. For the position controller, we use K_p and K_d gains of 40 and 5, respectively, and calculate torques for linearly interpolated joint angles from current to desired at 500Hz. We define the limits of the action space to be 30% of the physical joint limits.

TABLE VI: Simulated Reward Function Parameter Details

Parameter	Value
Target Velocity	1.0
Energy Penalty Weight	0.008
Qpos Penalty Weight	10.0
Smooth Torque Penalty Weight	0.005
Pitch Rate Penalty Factor	0.6
Roll Rate Penalty Factor	0.6
Joint Diagonal Penalty Weight	0.1
Joint Shoulder Penalty Weight	0.15
Smooth Change in Target Delta Yaw Steps	5

For the first experimental setting, we train prior behavior policies with high stiffness (10.0) in 9 different individual joints. Specifically, we use the front right body joint, the front right knee joint, the front left body joint, the front left knee joint, the rear right body joint, the rear right knee joint, the rear left body joint, the rear left thigh joint, and the rear left knee joint. During deployment, we switch between 3 conditions every 100 steps. Condition 1 is applying stiffness 15.0 to the rear right thigh joint, condition 2 is applying stiffness 15.0 to the front left thigh joint, and condition 3 is applying stiffness 15.0 to the front right thigh joint. For this setting, we use $\beta = 0.5$ for ROAM.

For the second experimental setting, we train prior behavior policies with low foot friction (0.4) in each of the 4 feet. During deployment, we switch between 2 conditions every 50 steps. Condition 1 is applying a foot friction of 0.1 to the rear right foot and condition 2 is applying a foot friction of 0.01 to the front left foot and a foot friction of 0.1 too the rear right foot. For this setting, we use $\beta = 0.5$ for ROAM.

b) Real-world Experiments.

For the real-world experiments, the state space consists of joint positions, joint velocities, torques, forward linear velocity, IMU (roll, pitch, change in roll, change in pitch), and normalized foot forces for a total of 47 dimensions. We use an Intel T265 camera-based velocity estimator to estimate onboard linear velocity. We use K_p and K_d gains of 20 and 1, respectively, which are used in the position controller. We again use action interpolation, an action range of 35% physical limits, and a 1 step action history. We also use a second-order Butterworth low-pass filter with a high-cut value of 8 to smooth the position targets. Finally, to reset the robot, we use the reset policy provided by smith2022walk. We train 4 prior behavior policies for the real-world experiments, each of which is trained with a frozen knee joint. Specifically, we train a policy with the front right knee joint frozen, the front left knee joint frozen, the rear right knee joint frozen, and the rear left knee joint frozen. $\beta = 0.5$ is used in all real-world experiments for ROAM.

TABLE VII: Real-world Reward Function Parameter Details

Parameter	Value
Target Velocity	1.5
Energy Penalty Weight	0.0
Qpos Penalty Weight	2.0
Smooth Torque Penalty Weight	0.005
Pitch Rate Penalty Factor	0.4
Roll Rate Penalty Factor	0.2
Joint Diagonal Penalty Weight	0.03
Joint Shoulder Penalty Weight	0.0
Smooth Change in Target Delta Yaw Steps	1

c) HLC Details.

For HLC, we have an MLP that takes state as input and outputs which behavior to select in the given state. The MLP has 3 hidden layers of 256 units each and ReLU activations, and we train by sampling from the combined offline data from all prior behaviors. We use a batch size of 256, learning rate of 3×10^{-4} , and train for 3,000 iterations.

d) RMA Details.

For RMA training, we changed the environment dynamics between each episode and trained for a total of 2,000,000 iterations. The standard architecture and hyperparameter choices from [19] were used along with DroQ hiraoka2021dropout as the base algorithm.