

# Task and Domain Adaptive Reinforcement Learning for Robot Control

Yu Tang Liu<sup>1,2</sup>, Nilaksh Singh<sup>3,2</sup>, and Aamir Ahmad<sup>2,1</sup>

**Abstract**—Deep reinforcement learning (DRL) has shown remarkable success in simulation domains, yet its application in designing robot controllers remains limited, due to its single-task orientation and insufficient adaptability to environmental changes. To overcome these limitations, we present a novel adaptive agent that leverages transfer learning techniques to dynamically adapt policy in response to different tasks and environmental conditions. The approach is validated through the blimp control challenge, where multitasking capabilities and environmental adaptability are essential. The agent is trained using a custom, highly parallelized simulator built on IsaacGym. We perform zero-shot transfer to fly the blimp in the real world to solve various tasks. We share our code at [https://github.com/robot-perception-group/adaptive\\_agent/](https://github.com/robot-perception-group/adaptive_agent/).

## I. INTRODUCTION

One promising approach to achieve optimal control in robotics is through reinforcement learning (RL) [1], an algorithm for searching control policies by environmental interaction. Despite its potential, RL methods remain impractical for most robotics applications, primarily due to the agents' inability to adapt to changes in key components of the Markov decision process (MDPs), such as reward function and transition dynamics.

In this work, we developed an adaptive agent that is robust against those changes or capable of adapting to those changes, by leveraging transfer learning (TL) techniques. Specifically, we focus on task transfer [2] and domain transfer [3]. Task transfer enhances the agent's ability to apply previously acquired skills to new tasks, thereby allowing building complex skills from simpler ones. Domain transfer, on the other hand, facilitates adaptation to varying transition dynamics, especially crucial for bridging the sim-to-real gap.

We introduce two fundamental modules developing our adaptive agent: (1) an architecture, which we call Arbiter-SF, that facilitates task transfer, and (2) a robust feature extractor. The task transfer module is inspired by the arbiter architecture [4], which involves an arbiter and multiple sub-policies, or "primitives," each specialized in handling specific sub-tasks. At every time step, these primitives propose their actions to the arbiter, which then evaluates these suggestions before the action selection. To enable task and domain transfer, we extend the arbiter observations with task specifications and the environment states (Fig. 1). To enhance task transfer performance, we integrate a *successor feature* (SFs, [5]) into the arbiter framework, which allows specifying a task as a

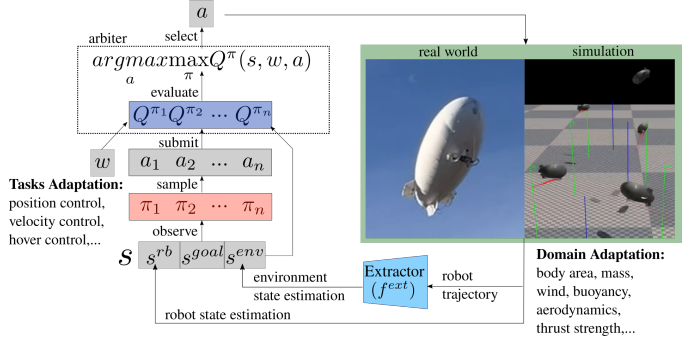


Fig. 1: The arbiter architecture allows our adaptive agent trained purely in simulation to achieve zero-shot domain transfer and sim-to-real transfer, to control the real robot and perform unseen tasks. The arbiter selects action by observing the current task weight  $w$ , state of the robot  $s^{rb}$ , goal  $s^{goal}$  and environment  $s^{env}$ . Our customized blimp simulation based on IsaacGym can support a high degree of parallelization to facilitate multitask learning. Each environment possesses a green and a blue pole corresponding to the hover and navigation goals' positions, respectively. The agent should navigate the blimp to different goal areas depending on the task specification.

continuous vector and enables real-time task transfer. Secondly, to adapt to varying dynamics, we infer environment states via a feature extractor by reading a snapshot of past interactions. To train this extractor, our approach adopts a robust two-stage training procedure, Rapid Motor Adaptation (RMA) [6], which trains an encoder first to capture the environmental factors within a domain randomized simulation and then replicates the representation by the extractor.

To validate our approach, we tested the adaptive agent in the autonomous blimp control challenge. Given a blimp's extended hover ability, they will undertake various control tasks, from basic maneuvers like hover, navigation, and take-off to activities such as aerial target tracking [7]. By learning these tasks together, which share a common state space, we can potentially accelerate learning and enhance the sample efficiency of the RL agent. Nevertheless, blimps' advantages come at the cost of complex dynamics: a highly nonlinear, time-delayed, environment-dependent system requiring policies that are either robust or adaptable to varying air conditions such as temperature, wind, and buoyancy. To address the computational demands of ROS/Gazebo simulation used in previous works [8], we developed a blimp simulator in IsaacGym [9], enabling significantly higher parallelization that facilitates multi-task training.

We summarize the contributions as follows:

- A novel adaptive RL agent via task and domain transfer that is tailored to the robotics applications. This is achieved through an Arbiter-SF architecture, a feature extractor, and an RMA training procedure.
- A customized blimp simulation, based on IsaacGym [9],

<sup>1</sup>Max Planck Institute for Intelligent Systems, 72076 Tübingen, Germany.  
<sup>2</sup>University of Stuttgart, 70569 Stuttgart, Germany. <sup>3</sup>IIT Kharagpur, West Bengal, India, 721302. [yutang.liu@tuebingen.mpg.de](mailto:yutang.liu@tuebingen.mpg.de), [nilaksh404@kqpian.iitkgp.ac.in](mailto:nilaksh404@kqpian.iitkgp.ac.in), [aamir.ahmad@ifr.uni-stuttgart.de](mailto:aamir.ahmad@ifr.uni-stuttgart.de)

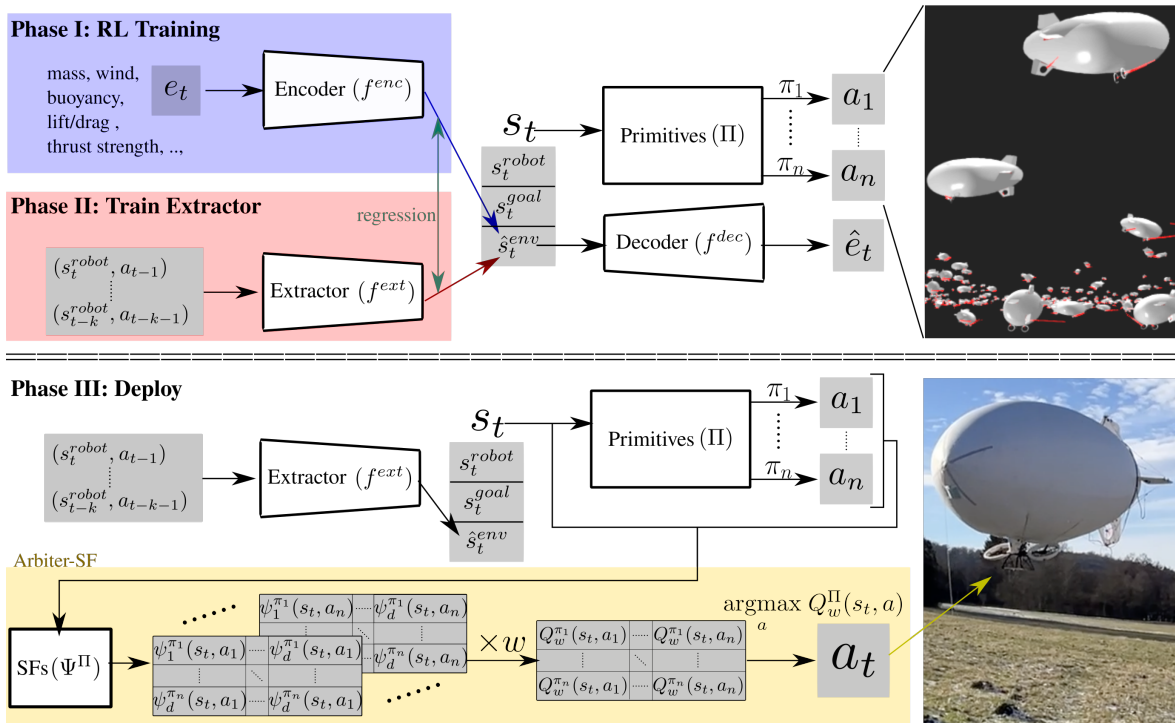


Fig. 2: **Our proposed methodology**: the adaptive agent follows the RMA training procedure. In the first phase, all modules except the feature extractor undergo training, focusing on training the primitives and constructing informative environmental states. In the second phase, the extractor is trained to replicate the encoder’s output, ensuring that it captures the environment state. A decoder is utilized in both phases to refine the environmental representation. During training, each primitive may control several blimps simultaneously. The arbiter is introduced during the deployment phase for selecting an action.

supporting thousands of parallel simulations.

## II. RELATED WORK

A direct approach to obtain a **multi-task agent** is augmenting task representations to the observation [10], which could demand extensive training data. An alternative is task transfer, where a trained policy is recycled to solve a new task such that agents can progressively handle harder challenges and learn a repertoire of skills [11]. There exist numerous task transfer methods. Primitive-based methods train a gating network for primitives composition [12], or chain in temporal sequences, e.g. hierarchical RL [13], options [14], etc. Modularized methods obtain a single multi-task policy via various sub-networks [15], or policy distilling [16]. Alternatively, representation learning develops a shared representation [17] or an explicit dynamics model [18] for solving new tasks. In contrast, our work adopts a unique kind of value function, i.e. SFs [19], facilitating *zero-shot* task transfer *without* augmenting task representation [20]. Our adaptive agent, built on top of the arbiter architecture [4] and SFs’ policy improvement theory [5], targets the *continuous control* domain and uses value functions for primitive composition, *without* training an additional gating network.

To bridge the **domain gaps** [3], domain randomization techniques enhance policy robustness by training across varied simulations [21], while domain adaptation aligns representations across source and target domains [22], often employing imitation learning [23]. Few-shot learning adapts to a new domain, either using gradient update [24], or estimate

environment state [25]. Our method is based on the RMA zero-shot approach to obtain a feature extractor that estimates environment state [6].

Prior work on **blimp control** uses ROS/Gazebo simulation and cascade PID controllers [26], encountering difficulties due to the blimps’ intricate dynamics. To handle environmental uncertainties like wind, buoyancy, and weight distribution, robust and adaptive control methods are desired. Learning-based approaches can potentially address these challenges [27]. A deep residual RL framework for position control has demonstrated superior performance by integrating classic methods [8]. In this work, we instead focus on enhancing multitasking capabilities and domain adaptability.

## III. METHODOLOGY

### A. Problem Definition

The goal of an adaptive agent is to search a set of domain adaptive primitives for task transfer via an arbiter. An adaptive agent improves a primitive through interaction with the environment. It formulates this procedure by MDPs  $\mathcal{M}_w \equiv \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, R_w, \gamma \rangle$ , where  $\mathcal{S}$  and  $\mathcal{A}$ , define the state and action space respectively. The dynamics  $\mathcal{T}(s_{t+1}|s_t, a_t) : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ , defines the transition probability from one state to another and  $\gamma \in [0, 1]$  is the discount factor. Each of the MDPs has a unique reward function  $r_t = R_w(s_t, a_t, s_{t+1}) : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$  specifies the desired task. We assume a linear reward function that can be decomposed as follows,  $R_w(s_t, a_t, s_{t+1}) = \Phi(s_t, a_t, s_{t+1})^\top \cdot w$ , where  $w \in \mathbb{R}^d$  is the

task weight and  $\phi_t = \Phi(s_t, a_t, s_{t+1}) \in \mathbb{R}^d$  is the features, which are relevant to task solving.

In task transfer, we consider a set of MDPs  $\{\mathcal{M}_w | w \in \mathcal{W}\}$ , and each member MDP possesses a unique task weight defined in the task space  $\mathcal{W}$  with the features shared with all other member MDPs. The goal is to find a set of  $n$  primitives,  $\Pi = \{\pi_{w_1}, \dots, \pi_{w_n}\}$ , where each primitive,  $\pi_{w_i}(a|s) : \mathcal{S} \rightarrow \mathcal{A}$ , is trained to solve its primitive task,  $w_i \in \mathcal{W}' \subset \mathcal{W}$ , by maximizing the expected discounted return,  $J(\pi_{w_i}) = \mathbb{E}_{\pi_{w_i}, \mathcal{M}_{w_i}} [\sum_{t=0}^{\infty} \gamma^t r_{w_i, t}]$ , such that we can recover the policies  $\pi_w$  of the unseen tasks  $w \in \mathcal{W}$ , by composing these primitives that only solves a subset of tasks  $\mathcal{W}'$ .

One way to transfer between domains is to extract the environment state from the observation. We model robot state, goal, together with the environment state, as part of the state space, i.e.  $S = (S^{robot}, S^{goal}, S^{env})$ , where  $s_t^{robot} \in S^{robot}$  can be observed via sensors or from a state estimator,  $s_t^{goal} \in S^{goal}$  is virtual goal states specified by user, and  $s_t^{env} \in S^{env}$  is the environment state un- (or partially) observable by the sensors. We aim to find a feature extractor,  $f^{ext}$  that can extract environment state  $s_t^{env}$  from a trajectory snapshot  $\tau_{t:t-k} = (s_t^{robot}, a_{t-1}, \dots, s_{t-k}^{robot}, a_{t-k-1})$ , i.e.  $\hat{s}_t^{env} = f^{ext}(\tau_{t:t-k})$ . Then the full state can be represented as  $s_t = (s_t^{robot}, s_t^{goal}, \hat{s}_t^{env})$ .

In summary, the task transfer problem boils down to training primitives  $\Pi = \{\pi_{w_1}, \dots, \pi_{w_n}\}$  for solving tasks  $\{w_1, \dots, w_n\}$  such that an arbiter can compose them for solving any tasks  $w$ . The domain transfer problem becomes training a feature extractor  $f^{ext}(\tau_{t:k})$  to reveal the environment state  $s_t^{env}$ .

### B. Task Transfer via Successor Feature-based Arbiter

When training the primitives, a set of value functions, or Q-functions,  $Q^\Pi = \{Q^{\pi_{w_1}}, \dots, Q^{\pi_{w_n}}\}$  are introduced. A Q-function is defined as:

$$Q_{\mathcal{M}_{w_i}}^{\pi_{w_i}}(s_t, a_t) \equiv \mathbb{E}_{\mathcal{M}_{w_i}, \pi_{w_i}} \left[ \sum_{k=0}^{\infty} \gamma^k r_{w_i, t+k} | s_t, a_t \right], \quad (1)$$

which estimates the total discounted reward of task  $w_i$  by following primitive  $\pi_{w_i}$ . One can see a value function as a performance report of a primitive, i.e., the higher the value indicating a better performance of the primitive in solving a task. Therefore, given a set of primitives, the arbiter is an operator,  $\operatorname{argmax}_{\pi_w \in \Pi} Q^{\pi_w}$ , reading values of all primitives for selecting the best one to perform the action. For brevity, we define  $Q_w^{\pi_{w_i}} \equiv Q_{\mathcal{M}_w}^{\pi_{w_i}}$ , meaning the performance of primitive  $\pi_{w_i}$  in solving task  $w$ . Given that  $R_{w_i}(s_t, a_t, s_{t+1}) = \Phi(s_t, a_t, s_{t+1})^\top \cdot w_i$ , we can rewrite the Q-function in following form,

$$Q_{w_i}^{\pi_{w_i}}(s_t, a_t) = \mathbb{E}_{\mathcal{M}_{w_i}, \pi_{w_i}} \left[ \sum_{k=0}^{\infty} \gamma^k r_{w_i, t+k} \right] \quad (2)$$

$$= \mathbb{E}_{\mathcal{M}_{w_i}, \pi_{w_i}} \left[ \sum_{k=0}^{\infty} \gamma^k \phi_{t+k}^\top \right] \cdot w_i = \psi^{\pi_{w_i}}(s_t, a_t)^\top \cdot w_i \quad (3)$$

where  $\psi^{\pi_{w_i}}(\cdot) \in \mathbb{R}^d$  is the *successor feature* (SFs) of the primitive  $\pi_{w_i}$ . Since Q-function can be viewed as a performance metric of a primitive, SFs provide a shortcut to estimate primitive performance in different tasks  $w$ , i.e.  $Q_w^{\pi_{w_i}}(s_t, a_t) = \psi^{\pi_{w_i}}(s_t, a_t)^\top \cdot w$ .

In the SF-based arbiter framework, training a number of primitives  $\Pi$  induces their associated SFs instead of Q-function, i.e.,  $\Psi^\Pi = \{\psi^{\pi_{w_1}}, \dots, \psi^{\pi_{w_n}}\}$ . Thus, depending on the current task  $w$ , the arbiter can select primitive actions based on profiles provided by the SFs [5], i.e.

$$\begin{aligned} \operatorname{argmax}_{a_t \sim \pi} \max_{Q_w^\pi \in Q^\Pi} Q_w^\pi(s_t, a_t) &\equiv \\ \operatorname{argmax}_{a_t \sim \Pi} Q_w^\Pi(s_t, a_t) &= \operatorname{argmax}_{a_t \sim \Pi} \Psi^\Pi(s_t, a_t)^\top \cdot w. \end{aligned} \quad (4)$$

which means to sample actions from each primitive and select the action incurring the largest primitive Q-values estimated via primitive SFs. Fig. 2 summarizes the action selection procedure.

### C. Successor Feature Policy Evaluation and Improvement

We can obtain primitives and their SFs via soft SFs policy iteration [20], which iterates between a soft SF policy evaluation step and a soft policy improvement step. The evaluation step improves the SFs estimation as follows,  $\psi^{\pi_w}(s_t, a_t)$

$$= \phi_t(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim p} \left[ \alpha \log \int_{a \sim \pi_w} e^{\psi^{\pi_w}(s_{t+1}, a)/\alpha} \right], \quad (5)$$

where  $\alpha \in \mathbb{R}^d$  is a temperature parameter for controlling the exploration of each primitive. Then the soft policy improvement step aims to replicate policy distribution based on the Q-values of each action, i.e., the larger the Q-value, the higher the probability of sampling the action,  $\pi_w(\cdot | s) \propto e^{Q_w^{\pi_w}(s, \cdot)}$ , which is achieved by solving,

$$\pi_w(\cdot | s) = \operatorname{argmin}_{\pi'_w} D \left( \pi'_w(\cdot | s_t) \left\| \frac{e^{Q_w^{\pi_w}(s_t, \cdot)}}{Z(s_t)} \right. \right), \quad (6)$$

where  $Q_w^{\pi_w}$  is derived by (3),  $D(\cdot | \cdot)$  measures the KL divergence, and  $Z(s_t) = \int_{\mathcal{A}} e^{Q_w^{\pi_w}(s_t, \cdot)}$  is a negligible function as it does not contribute to gradient computation. Alternating between an evaluation step and an improvement step guarantees convergence to the optimal primitive  $\pi_{w_i}^*$  and the optimal Q-function  $Q_{w_i}^{*, \pi_{w_i}} \equiv Q_{w_i}^{\pi_{w_i}^*}$  yielded from the optimal SFs  $\psi^{*, \pi_{w_i}} \equiv \psi^{\pi_{w_i}^*}$  of each primitive.

### D. Training Scheme

The RMA [6] unfolds training in three phases (Fig. 2): Phase I: Training the primitive module  $\Pi$ , SFs module  $\Psi^\Pi$  and environment encoder  $f^{enc}$ . Phase II: Replicating the encoder  $f^{enc}$  representation using the feature extractor  $f^{ext}$ . Phase III: Evaluation/deployment phase.

1) *Phase I: Training the SFs and Primitive Network*: The primitives are trained by alternating between an interaction phase and a training phase. During the interaction phase, training data is gathered by interaction within the environment in parallel and stored in a replay buffer. In the training phase, we sample data from the replay buffer and update the SFs and primitives via (5) and (6), respectively.

Every training episode, the simulation exposes randomized environment factors of length  $E$ , i.e.,  $e \in \mathbb{R}^E$ . Then the encoder, parameterized by  $\theta_{f^{enc}}$ , learns to convert these factors to the environment state,  $\hat{s}_t^{env} = f^{enc}(e_t)$ . Then the state variable is a concatenation of the robot state, goal state, and

the environment state  $s_t = (s_t^{robot}, s_t^{goal}, \hat{s}_t^{env})$ , where robot states and goal states are randomized for each training episode.

In total, we have  $n$  pairs of SFs and primitives with the  $i$ -th successor feature  $\psi^{\pi_{w_i}}$ , target network  $\hat{\psi}^{\pi_{w_i}}$ , and the primitive network  $\pi_{w_i}$  parameterized by  $\theta_{\psi^{\pi_{w_i}}}$ ,  $\theta_{\hat{\psi}^{\pi_{w_i}}}$ , and  $\theta_{\pi_{w_i}}$  respectively. The target network is a common practice in RL to reduce the overestimation of the values. The loss for the  $i$ -th successor feature is computed as follows,

$$J_{\psi^{\pi_{w_i}}}(\theta_{\psi^{\pi_{w_i}}}) = \mathbb{E}_{(s_t, a_t) \sim \mathcal{D}} \left[ \|\psi^{\pi_{w_i}}(s_t, a_t) - \hat{\psi}^{\pi_{w_i}}(s_t, a_t)\|_2 \right], \quad (7)$$

where  $\mathcal{D}$  is the replay buffer and  $\|\cdot\|_2$  denotes L2 norm. We have the TD-target  $\hat{\psi}^{\pi_{w_i}}$  approximated by the target network  $\bar{\psi}^{\pi_{w_i}}$ ,

$$\hat{\psi}^{\pi_{w_i}}(s_t, a_t) = \phi_t + \gamma \mathbb{E}_{\substack{s_{t+1} \sim \mathcal{D} \\ a_{t+1} \sim \pi_{w_i}}} [\bar{\psi}^{\pi_{w_i}}(s_{t+1}, a_{t+1})]. \quad (8)$$

Then the primitive can be improved by minimizing the objective,

$$J_{\pi_{w_i}}(\theta_{\pi_{w_i}}) = -\mathbb{E}_{\substack{s_t \sim \mathcal{D} \\ a_t \sim \pi_{w_i}}} [Q_{w_i}^{\pi_{w_i}}(s_t, a_t) - \alpha \log \pi_{w_i}(a_t | s_t)], \quad (9)$$

where  $Q_{w_i}^{\pi_{w_i}}(s_t, a_t) = \psi^{\pi_{w_i}}(s_t, a_t)^\top \cdot w_i$  and  $-\log \pi(\cdot | s)$  is an estimator for the policy entropy.

Note that, in theory, we can update the encoder  $\theta_{f^{enc}}$  via policy and SFs loss. However, in practice, we found training much more stable if only the SFs loss is applied.

2) *Phase II: Training the Feature Extractor*: We train a feature extractor parameterized by  $\theta_{f^{ext}}$  to replicate the environment state predicted by the encoder by minimizing the loss,

$$J_{f^{ext}}(\theta_{f^{ext}}) = \|f^{enc}(e_t) - f^{ext}(\tau_{t:t-k})\|_2. \quad (10)$$

We observed instability in training when initiated with only the feature extractor (Fig.6a). Therefore, training an encoder first becomes necessary.

Lastly, during deployment (Phase III), we use arbiter for composing primitive actions using (4). Algorithm 1 summarizes how the adaptive agent trains each network module. Note that, to keep the computation time constant, in practice, the *for*-loop is replaced with vectorized computations.

### E. Supporting Techniques

We introduce compatible techniques for the adaptive agent to improve training stability and sample efficiency.

1) *Collective Learning*: An alternative to (6) is to use the collective knowledge of multiple SFs. One can see all SFs,  $\Psi^\Pi = \{\psi^{\pi_{w_1}}, \dots, \psi^{\pi_{w_n}}\}$ , as an ensemble bundled to produce improved value estimation. In theory, the Q-value incurred from the ensemble,  $Q_w^{sup} = \max_{\psi^\pi \in \Psi} \psi^\pi \cdot w$ , is guaranteed to be greater or equal to any primitive Q-values in any task  $w$  [5], i.e.  $Q_w^{sup} \geq \sup_{\pi \in \Pi} Q_w^\pi$ . That is, we can replace  $Q_{w_i}^{\pi_{w_i}}$  with  $Q_{w_i}^{sup}$  to conduct policy improvement step (6) for  $\pi_{w_i}$ .

2) *Imitation Learning*: The arbiter architecture allows injecting prior knowledge to primitives via behavior cloning (BC) by minimizing the discrepancy between student and teacher actions:

$$J_{BC}(\theta_{\pi_{w_i}}) = \mathbb{E}_{s \sim D} \mathbb{E}_{\substack{a^{\pi_{w_i}} \sim \pi_{w_i} \\ a^{\pi_{tea}} \sim \pi_{tea}}} [\|a^{\pi_{w_i}} - a^{\pi_{tea}}\|_2], \quad (11)$$

where  $\pi_{tea}$  is a teacher policy. We need to assign primitives with their corresponding teachers and increment the BC loss

### Algorithm 1 Adaptive RL Agent

---

```

Initialize network parameters  $\theta_{\psi^\pi}, \theta_{\hat{\psi}^\pi}, \theta_\pi, \theta_{f^{enc}}, \theta_{f^{ext}}$ 
1: while Training do
2:   if Phase==1 then  $\hat{s}_t^{env} = f^{enc}(e_t)$ 
3:   else  $\hat{s}_t^{env} = f^{ext}(\tau_{t:t-k})$ 
4:    $s_t = (s_t^{robot}, s_t^{goal}, \hat{s}_t^{env})$ 
5:
6:   if exploration then  $a_t \leftarrow \text{Uniform}(A)$ 
7:   else if Phase==1 or 2 then
8:      $a_t \leftarrow \pi_i(a|s_t), w_i \sim \mathcal{W}'$ 
9:   else if Phase==3 then
10:     $a_t \leftarrow \arg\max_{a_t \sim \Pi} \Psi^\Pi(s_t, a_t)^\top \cdot w, w \sim \mathcal{W}$ 
11:   $\mathcal{D} \leftarrow \mathcal{D} \cup (s_t, a_t, s_{t+1}, \Phi(s_t, a_t, s_{t+1}))$ 
12:
13:  if Phase==1 then
14:    for  $i \leftarrow 1, 2, \dots, n$  do
15:       $\theta_{\psi^{\pi_{w_i}}} \leftarrow \theta_{\psi^{\pi_{w_i}}} - \lambda_\psi \nabla_{\theta_{\psi^{\pi_{w_i}}}} J_{\psi^{\pi_{w_i}}}(\theta_{\psi^{\pi_{w_i}}})$ 
16:       $\theta_{\pi_{w_i}} \leftarrow \theta_{\pi_{w_i}} - \lambda_\pi \nabla_{\theta_{\pi_{w_i}}} J_{\pi_{w_i}}(\theta_{\pi_{w_i}})$ 
17:       $\theta_{\hat{\psi}^{\pi_{w_i}}} \leftarrow \tau \theta_{\hat{\psi}^{\pi_{w_i}}} + (1 - \tau) \theta_{\psi^{\pi_{w_i}}}$ 
18:    else if Phase==2 then
19:       $\theta_{f^{ext}} \leftarrow \theta_{f^{ext}} - \lambda_f \nabla_{\theta_{f^{ext}}} J(\theta_{f^{ext}})$ 
20:    else if Phase==3 then pass

```

---

to the respective primitive losses (9), i.e.  $J_{\pi_{w_i}}(\theta_{\pi_{w_i}}) += k_{BC} \cdot J_{BC}(\theta_{\pi_{w_i}})$ , where  $k_{BC} \in \mathbb{R}^+$ . This way, those primitives will replicate teacher actions and improve upon them based on their own primitive tasks.

3) *Auxiliary Task*: The auxiliary task can improve the sample efficiency and representation by making auxiliary predictions [28] or solving auxiliary control task [29]. In our case, one SFs head  $j$ , where  $j > n$ , is assigned to predict the next features for increasing the dynamics awareness by minimizing the objective,

$$J_{Aux}(\theta_{\psi^{aux_j}}) = \mathbb{E}_{\substack{s_t, a_t \sim D \\ s_{t+1}, a_{t+1} \sim D}} [\|\psi^{aux_j}(s_t, a_t) - \phi_{t+1}\|_2]. \quad (12)$$

where  $k_{Aux} \in \mathbb{R}^+$ . Additional auxiliary control tasks are introduced in Sec. IV-A.2.

4) *Specialized Network Architecture*: The goal of network architecture (Fig. 2) is to develop a representation that can be shared among the primitives and avoid conflicting gradients [30] to stabilize multitask learning. One way to reduce the conflict is by increasing the sparsity, for which we embed a fuzzy tiling activation function (FTA) before the final layer of primitive network [31]. We adopt residual network architecture [32] and decoder  $f^{dec}$  to construct the deeper network to preserve the information on environmental factors and enhance the agent's performance. The decoder loss is described as follows,

$$J_{f^{dec}}(\theta_{f^{dec}}, \theta_{f^{enc}}) = \|e_t - f^{dec}(f^{enc}(e_t))\|_2. \quad (13)$$

It is incremented with SFs loss (7) to update the encoder.

## IV. EXPERIMENTS AND RESULTS

The experiments aim to answer the questions: (1) How well is the adaptive agent solving unseen tasks? (2) What are the critical aspects of successfully training an adaptive agent?

### A. Blimp Control Problem

1) *State, Action, and Feature Space*: The blimp control's state, action, and feature space are described in Table. 1. At every time step, each primitive observes the state of the




Fig. 3: **Tasks performed in the simulation:** Various kind of tasks can be achieved by mixing components in the task weight. (a) The hover task  $w[4]$  can be combined with yaw control task  $w[5]$  to track desired angle in the hover zone  $w = [0, 0, 0, 0, 1, 1, 0, 0, 0, 0]$ . We can also add a thrust penalty ( $w[10] = 0.1$ ) to prevent overshooting (brown). (b) The hover task  $w[4]$  can be combined with velocity control  $w[6]$  to track desired velocity in the hover zone  $w = [0, 0, 0, 0, 1, 0, 1, 0, 0, 0]$ . We can reduce the velocity weight to emphasize less on this sub-task (brown), i.e.,  $w_v = w[6] = 0.1$ . (c) A higher control performance in waypoint trigger task  $w[2]$  can be achieved by mixing position  $w[0]$  and velocity control tasks  $w[7]$ , i.e.  $w = [.3, .1, 1, 0, 0, 0, 0, .3, 0, 0, 0]$ . Compared with a velocity PID controller (brown), the agent does not blindly follow the velocity commands to prevent overshooting. (d) The agent can fly the blimp backward  $w = [0, 0, 0, -1, 0, 0, -1, 0, .1, 0, 0]$  by reversing the yaw  $w[3]$  and heading velocity task  $w[6]$ . However, backward flight is slow and inefficient. The agent can hardly maintain its altitude and ends up being reset by the simulator.




Fig. 4: **Experimental results on task transfer:** The solid line indicates the mean, and the top and bottom edges of the hue are the maximum and minimum rewards of each experiment. Compared to (a), the learning rate in experiments (b) and (c) is reduced ten times, but the number of episodes is increased from 125 to 200 to ensure training stability for better evaluation of the supporting techniques. Note that these experiments are conducted without domain randomization. Environment factors are set to default values and excluded from the agent’s observation. (a) Benchmark on transfer performance in the evaluation task set. The green dotted line indicates the final performance of the agent deployed for real-world experiments. All agents incorporate imitation loss (Sec. III-E.2). (b) All the techniques, except collective learning, significantly improve the sample efficiency and transfer performance. *base*: no techniques, *imi*: imitation learning, *fta*: fuzzy tiling activation, *aux*: predictive auxiliary task, *col*: collective learning. (c) Training the Arbiter-SF with more tasks may improve final performance (115 for  $N = 10$  and 80 for  $N = 5$ ) at the cost of learning time.

robot, goal, and environment and produces an action to control the blimp. The state space is constituted of the robot, goal,

$\mathcal{S}$	$pose^{robot}, pose^{goal_{nav}}, pose^{goal_{hov}}, s^{env}$
$\mathcal{A}$	$a^{thrust}, a^{servo}, a^{pitch}, a^{yaw}$

TABLE 1: State and action space for the blimp control

and environment states, where  $pose = (\mathbb{X}, \mathbb{V}, \Theta, \Omega) \in (\mathbb{R}^3)^4$  denotes the position, velocity, angle, and angular velocity of the robot and goals. For convenience, we define two goals, a navigation and a hover goal (Fig. 1). The hover goal  $pose^{goal_{hov}}$  is placed at the center of the flight zone, whereas the navigation goals  $pose^{goal_{nav}}$  are waypoints defined within the flight zone. The extracted environment state  $s^{env} \in \mathbb{R}^{15}$  has an entry length defined by the user. In practice, actuator states are included, e.g., thrust and servo angle. The action space is a vector of four entries  $\mathcal{A} \in [-1, 1]^4$  (Table. 1), where  $a^{thrust}$  controls two motors,  $a^{servo}$  controls the thrusting angle,  $a^{pitch}$  controls the left and the right elevators, and  $a^{yaw}$  controls the top and bottom rudders and a bottom motor.

We define the task-relevant features  $\phi \in \mathbb{R}^{11}$  as follows (14),

$$\phi = \begin{bmatrix} \phi^{dist_{xy}} \\ \phi^{dist_z} \\ \phi^{trigger} \\ \phi^{yaw_{heading}} \\ \phi^{proximity} \\ \phi^{yaw} \\ \phi^{v_{heading}} \\ \phi^{v_{xy}} \\ \phi^{v_z} \\ \phi^{reg_{RP}} \\ \phi^{reg_T} \end{bmatrix} = \begin{bmatrix} \|\mathbb{X}_{xy}^{robot} - \mathbb{X}_{xy}^{goal_{nav}}\|_{\mathcal{N}} \\ \|\mathbb{X}_z^{robot} - \mathbb{X}_z^{goal_{nav}}\|_{\mathcal{N}} \\ 1 \text{ if } \|\mathbb{X}^{robot} - \mathbb{X}^{goal_{nav}}\|_2 < 5 \text{ else } 0 \\ \|\Theta_z^{robot} - \Theta_z^{goal_{nav}}\|_{\mathcal{N}} \\ 1 \text{ if } \|\mathbb{X}^{robot} - \mathbb{X}^{goal_{hov}}\|_2 < 7 \text{ else } 0 \\ \|\Theta_z^{robot} - \Theta_z^{goal_{hov}}\|_{\mathcal{N}} \\ \|\mathbb{v}_{\mathcal{B}}^{robot} - \mathbb{v}^{goal_{hov}}\|_{\mathcal{N}} \\ \|\mathbb{v}_{xy}^{robot} - \mathbb{v}_{xy}^{goal_{nav}}\|_{\mathcal{N}} \\ \|\mathbb{v}_z^{robot} - \mathbb{v}_z^{goal_{nav}}\|_{\mathcal{N}} \\ \|\Theta_{xy}^{robot}\|_{\mathcal{N}} \\ \|\mathbb{a}_{thrust} + 1\|_2 \end{bmatrix} \quad (14)$$

where  $\|\cdot\|_2$  denotes L2 norm, the Gaussian norm is defined as  $\|\cdot\|_{\mathcal{N}} \equiv e^{-k\|\cdot\|_2/\sigma^2}$  and subscript  $\mathcal{B}$  denotes state in body frame. These features measure the differences between robot and goal states for defining tasks. The first four are relevant to position control, the next three to hover control, the eighth and ninth to velocity control, and the last to regularization tasks.

2) *Tasks:* We have in total ten primitive tasks as the training task set  $w = \{w_1, \dots, w_{10}\}$  and each has the same dimension as the features, i.e.,  $w_i \in \mathbb{R}^{11}$ , defined in (15). The first three tasks of interest are position control, hover control, and velocity




Fig. 5: The color coding indicates the counts of each primitive being chosen to perform the action by the arbiter in an evaluation episode every five training episodes.




Fig. 6: **Experiment results on domain transfer:** The label indicates if the decoder/residual network is applied, and T or F denotes True or False. The experiments apply domain randomization and utilize the best configuration from 4b with an encoder and extractor incorporated. (a) The yellow line indicates learning with only the extractor. All configurations drop in performance when entering the second phase (Green). (b) The extractor loss (10) is higher with the decoder but reduced with the residual network.

control, denoted by  $w_1$ ,  $w_2$ , and  $w_3$ , respectively. In position control, the objective is to sequentially trigger waypoints by reducing the distance to the current one to be less than five meters. Then, the next waypoint will become activated. The hover task requires the blimp to maintain its position within a specified range for the longest possible duration. Velocity control, similar to position control, involves activating waypoints, but it is guided by velocity commands rather than positional ones. A planner provides velocity commands based on the relative position to navigation goal at every time step, i.e.  $\nabla^{goal_{nav}} = 6 \cdot (\mathbb{X}^{goal_{nav}} - \mathbb{X}^{robot}) / \|\mathbb{X}^{goal_{nav}} - \mathbb{X}^{robot}\|_1$ . The remaining tasks  $w_4 - w_{10}$  serve as auxiliary control tasks (as mentioned in Sec. III-E.3), offering alternative methods for tackling similar tasks or enhancing specific abilities, e.g.,  $w_8$  focuses solely on yaw angle control.

$$\mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \\ w_5 \\ w_6 \\ w_7 \\ w_8 \\ w_9 \\ w_{10} \end{bmatrix} = \begin{bmatrix} .1, .1, 1, 0, & 0, 0, 0, & 0, 0, & .01, 0 \\ 0, 0, 0, 0, & 10, 0, .01, & 0, 0, & 0, .1 \\ 0, 0, 0, 0, & 0, 0, 0, & 1, .5, & .01, 0 \\ .1, .1, 0, -1, & 0, 0, -1, & 0, 0, & 0, 0 \\ .1, .5, .1, 1, & 0, 0, 1, & 0, 0, & .01, .01 \\ 0, 0, 0, 0, & 10, .2, 0, & 0, 0, & 0, .5 \\ 0, 0, .1, .1, & 0, 0, 0, & 1, .5, & .01, 0 \\ 0, 0, 0, 0, & 0, 1, 0, & 0, 0, & 0, .1 \\ 1, .1, 0, .1, & 0, 0, 0, & 0, 0, & 0, 0 \\ 0, 0, 0, 0, & 1, 0, 0, & 0, 0, & .1, .1 \end{bmatrix} \cdot (15)$$

To accelerate training, knowledge from PID control-based teachers is incorporated through behavior cloning (Sec. III-E.2) to assist in solving the first four tasks  $w_1 - w_4$ . Given that PID teachers are optimized for nominal environment conditions and perform sub-optimally outside these conditions, blindly imitating them can lead to poor performance for the primitives.

$$\mathbf{w}^{eval} = \begin{bmatrix} w_1^{eval} \\ w_2^{eval} \\ w_3^{eval} \\ w_4^{eval} \\ w_5^{eval} \\ w_6^{eval} \\ w_7^{eval} \end{bmatrix} = \begin{bmatrix} 1, 0, 0, 0, & 0, 0, 0, & 0, 0, & 0, 0 \\ 0, 0, 1, 0, & 0, 0, 0, & 0, 0, & 0, 0 \\ 0, 0, 0, 1, & 0, 0, 1, & 0, 0, & 0, 0 \\ 0, 0, 0, 0, & 1, 0, 0, & 0, 0, & 0, 0 \\ 0, 0, 0, 0, & 1, 1, 0, & 0, 0, & 0, 0 \\ 0, 0, 0, 0, & 0, 0, 0, & 1, 0, & 0, 0 \\ 0, 0, 0, 0, & 0, 0, 0, & 1, .5, & 0, 0 \end{bmatrix} \cdot (16)$$

For evaluation, we define a set of evaluation tasks  $\mathbf{w}^{eval}$  (16), different from the training set, to test the agent’s transfer ability. These tasks also cover basic tasks, including position, velocity, and hover control, with some arbitrarily defined new unseen tasks. For each of the parallelized simulations, we

sample a task uniform randomly from the task set and apply an arbiter to perform task transfer.

### B. Experiment Setup

We build our simulation based on IsaacGym [9] to facilitate data collection in parallel and enable multi-task learning. The adaptive agent is built on Pytorch with just-in-time compilation to optimize the code performance.

For the blimp control, the first RMA training phase is set to 1000 episodes (13.9 hrs) and 250 episodes (3.8 hrs) for the second phase, with both using the training task set (15). Each episode contains 500 environment steps with 1024 environments running in parallel, with each environment step set to 0.1 seconds. The number of episodes is reduced in the ablation study to save the amount of computation. In every episode, domain randomization is applied to generate goals with different poses and environmental factors, as shown in Table. 2. To ensure the motion smoothness, instead of directly mapping the agent action to the thrust, a motor smoothing factor is applied via moving average with previous actions, i.e.,  $thrust_t = c_1 \cdot (c_2 \cdot a_t^{thrust} + (1 - c_2) \cdot a_{t-1}^{thrust})$ , where  $(c_1, c_2) = (10, 0.2)$ .

Type A	body area, weight distribution, wind magnitude and variance, blimp mass, lift/drag coefficients, buoyancy
Type B	thrust strength ( $c_1$ ), motor smoothing factor ( $c_2$ )

TABLE 2: Environment Factors. Type A variables are randomized in the range of 80-125% from the nominal condition, whereas 50-200% for type B variables. Note that the mean is set to the mass and drawn from the uniform distribution to keep buoyancy close to the mass.

The real-world experiment is performed on a five-meter blimp (Fig. 1), equipped with a GPS sensor and a Librepilot onboard flight controller, which has a built-in IMU and a gyroscope sensor. The state estimation is then obtained via an extended Kalman filter. The adaptive agent is mounted on the onboard GPU, Nvidia TX2, and runs at 10Hz. The task is assigned by the user on the laptop and transmitted to the agent via ROS multi-master service in *real-time*.

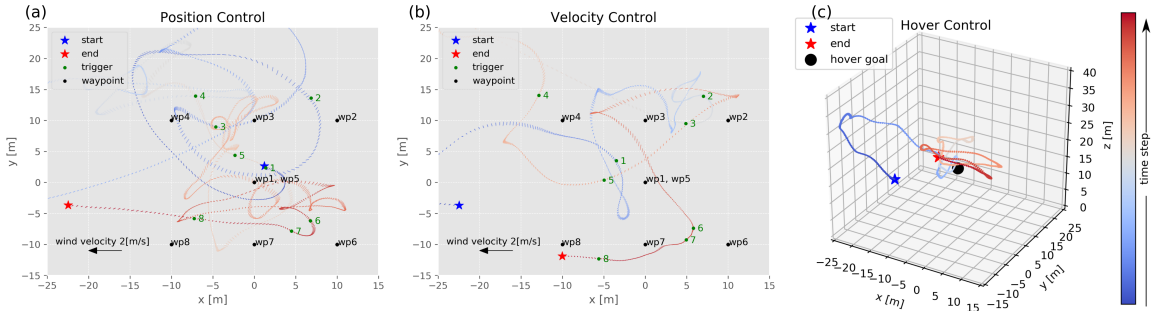


Fig. 7: **Real world experiments:** robot trajectories for different tasks performed in the real-world experiment. For visual clarity, only tasks involving changes in altitude are depicted in 3-D format. (a) The agent triggers all the waypoints by following the position command.  $w = [.5, .2, 1, .1, 0, 0, 0, 0, 0, .1, 0]$ . (b) The agent triggers all the waypoints by following the velocity command.  $w = [0, 0, 1, 0, 0, 0, 0, 1, .5, .1, 0]$ . (c) Hover tasks bring the blimp near the hover goal at the center of flight zone  $w = [0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 1]$ .

### C. Experiment Result

1) **Task Transfer in Simulated Environment:** In nominal environment conditions, our agent can perform zero-shot transfer to arbitrary tasks, as depicted in Fig. 3. The hover and angle control task (Fig. 3a) positions the blimp in the hover region with the desired angle. The task is useful for tracking with a fixed camera [7]. By increasing the weight of thrust regularization, we can prevent overshooting and reduce energy consumption. Integrating hover and velocity control task (Fig. 3b) allows to track the desired velocity within hover region. Sometimes, it is more efficient to stay in motion by leveraging aerodynamics over constant vertical thrusting. The position and velocity control (Fig. 3c) can be combined to improve waypoint tracking. Blindly following the velocity command often leads to overshooting, like the velocity PID teacher, and including position control can mitigate this. Lastly, the backward position control (Fig. 3d) cannot be achieved. The agent learns to point the blimp’s tail to the waypoint without maintaining a constant velocity. Still, backward control is practical for reducing velocity or when only a tiny correction from overshooting is required.

2) **Asymptotic Performance and Task Transfer-ability:** We benchmark our adaptive agent – SF-based arbiter (Arbiter-SF) – against other SAC-based agents for the blimp control task on asymptotic task transfer performance:

- SAC [33]: A state-of-the-art single task agent. Task weight is augmented to the observation for both Q-value and policy networks, i.e.,  $s \leftarrow (s, w)$ .
- Arbiter-Q: A Q-function-based arbiter derived from SAC. Task weight is augmented to the observation for both Q-value and primitive networks.
- Soft-Module [15]: A state-of-the-art multi-task agent with task embedding replaced by the task weight.

All agents are trained with the task set  $w$  (15) and evaluated with  $w^{eval}$  (16). To ensure a fair comparison, all supporting techniques were excluded from the Arbiter-SF, except for the imitation loss, which was incorporated into all agents, i.e., Arbiter-Q, SAC, and Soft-Module, to promote quicker convergence. Following a hyper-parameter tuning by the Bayesian optimization 25 times, the top performance for each agent was selected for comparison. For the blimp control, our

Arbiter-SF achieves the best transfer performance after 125 episodes of training (Fig. 4a).

3) **Ablation Study on Task Transfer:** All techniques (Sec. III-E) contribute significantly except collective learning (Fig. 4b), implying that the estimated Q-values are nearly identical by using the primitive’s SFs or by the collective SFs. Next, Fig. 4c shows that there is a trade-off between the number of primitives on learning speed and final performance, contrary to the idea that incorporating auxiliary tasks should increase sample efficiency [29], [28].

4) **Primitive Evaluation:** The learning pace of the primitives varies significantly, as shown in Fig. 5. At the start of training, the agent predominantly activates the first four primitives, which imitate PID teachers. Notably, the fourth primitive, guided by a backward flight PID, has the highest activation frequency. This may indicate that one initial learning outcome is decelerating the blimp’s speed. As training progresses, the activation frequency of primitives becomes more balanced, indicating that most primitives have developed skills in task-solving.

5) **Ablation Study on Domain Transfer:** Fig. 6a shows that applying residual networks for encoder and extractor improves transfer performance, but adding a decoder decreases it. Further examination in the second phase (Fig. 6b) reveals that the presence of a decoder results in significantly higher loss, suggesting that the resulting representation is more complex and potentially richer in content. However, this increased richness does not translate to more effective task-solving content with agents trained in limited episodes. Still, we recommend incorporating a decoder into the architecture as it can benefit agents trained over longer episodes.

6) **Real World Experiment:** The trajectory of the blimp is displayed in Fig. 7. Initially, a human pilot manually takes the blimp off to around 7 meters for safety purposes, after which the agent takes over the control to complete the ascent to 20 meters within the designated hover region. Then, the tasks, including position, velocity, and hover control, are executed sequentially. In each starting phase of the task, the human operator needs to tune the task weight. During the experiment, we observed a different flight pattern from the classic PID controller (Fig. 8). PID controllers generally rely on forward or upward thrust to maneuver the blimp, struggling with

thrust angles between or beyond ninety degrees due to non-linearity and the dynamic inversion problem. In contrast, the agent extensively employs thrust vectoring, achieving a more agile flight style. We refer more discussion on DRL and PID methods to [8], where our PID teachers were benchmarked.




Fig. 8: Notable behaviors include extensive use of the thrust vector (yellow arrow), such as forward (left), upward (middle), and backward thrust (right).

## V. CONCLUSIONS AND FUTURE WORK

In this work, we proposed an adaptive agent that supports task and domain transfer utilizing the arbiter architecture with value estimation provided by the successor feature. The experiments on the blimp control tasks demonstrate that our agent achieves the best transfer performance on unseen tasks and enables zero-shot transfer to different tasks and domains, validated in real-world experiments. One drawback, as mentioned in Sec. IV-C.3, is that providing more auxiliary control tasks did not always improve sample efficiency. Furthermore, sometimes, we observe action inconsistency, which implies that the SFs' value estimation can be noisy and disturb the arbiter's decision-making. Our future work will focus on increasing the action consistency and examining the side effects accompanied by the auxiliary control tasks.

## REFERENCES

- [1] D. Bertsekas and J. N. Tsitsiklis, *Neuro-dynamic programming*. Athena Scientific, 1996.
- [2] Z. Zhu, K. Lin, A. K. Jain, and J. Zhou, "Transfer learning in deep reinforcement learning: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [3] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: a survey," in *2020 IEEE symposium series on computational intelligence (SSCI)*. IEEE, 2020, pp. 737–744.
- [4] S. J. Russell and A. Zimdars, "Q-decomposition for reinforcement learning agents," in *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, 2003, pp. 656–663.
- [5] A. Barreto, S. Hou, D. Borsa, D. Silver, and D. Precup, "Fast reinforcement learning with generalized policy updates," *Proceedings of the National Academy of Sciences*, vol. 117, no. 48, pp. 30079–30087, 2020.
- [6] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "Rma: Rapid motor adaptation for legged robots," *arXiv preprint arXiv:2107.04034*, 2021.
- [7] E. Price, M. J. Black, and A. Ahmad, "Viewpoint-driven formation control of airships for cooperative target tracking," *IEEE Robotics and Automation Letters*, 2023.
- [8] Y. T. Liu, E. Price, M. J. Black, and A. Ahmad, "Deep residual reinforcement learning based autonomous blimp control," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 12 566–12 573.
- [9] J. Liang, V. Makovychuk, A. Handa, N. Chentanez, M. Macklin, and D. Fox, "Gpu-accelerated robotic simulation for distributed reinforcement learning," in *Conference on Robot Learning*. PMLR, 2018, pp. 270–282.
- [10] L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley, I. Dunning, *et al.*, "Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures," in *International conference on machine learning*. PMLR, 2018, pp. 1407–1416.
- [11] D. Kalashnikov, J. Varley, Y. Chebotar, B. Swanson, R. Joschkowski, C. Finn, S. Levine, and K. Hausman, "Mt-opt: Continuous multi-task robotic reinforcement learning at scale," *arXiv preprint arXiv:2104.08212*, 2021.
- [12] 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.
- [13] S. Pateria, B. Subagdja, A.-h. Tan, and C. Quek, "Hierarchical reinforcement learning: A comprehensive survey," *ACM Computing Surveys (CSUR)*, vol. 54, no. 5, pp. 1–35, 2021.
- [14] R. S. Sutton, D. Precup, and S. Singh, "Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning," *Artificial intelligence*, vol. 112, no. 1-2, pp. 181–211, 1999.
- [15] R. Yang, H. Xu, Y. Wu, and X. Wang, "Multi-task reinforcement learning with soft modularization," *Advances in Neural Information Processing Systems*, vol. 33, pp. 4767–4777, 2020.
- [16] D. Ghosh, A. Singh, A. Rajeswaran, V. Kumar, and S. Levine, "Divide-and-conquer reinforcement learning," *arXiv preprint arXiv:1711.09874*, 2017.
- [17] A. Gupta, C. Devin, Y. Liu, P. Abbeel, and S. Levine, "Learning invariant feature spaces to transfer skills with reinforcement learning," *arXiv preprint arXiv:1703.02949*, 2017.
- [18] A. Zhang, H. Satija, and J. Pineau, "Decoupling dynamics and reward for transfer learning," *arXiv preprint arXiv:1804.10689*, 2018.
- [19] P. Dayan, "Improving generalization for temporal difference learning: The successor representation," *Neural computation*, vol. 5, no. 4, pp. 613–624, 1993.
- [20] Y. T. Liu and A. Ahmad, "Multi-task reinforcement learning in continuous control with successor feature-based concurrent composition," in *2024 European Control Conference (ECC)*. IEEE, 2024, pp. 3860–3867.
- [21] 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.
- [22] M. E. Taylor and P. Stone, "Transfer learning for reinforcement learning domains: A survey," *Journal of Machine Learning Research*, vol. 10, no. 7, 2009.
- [23] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in neural information processing systems*, vol. 29, 2016.
- [24] 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.
- [25] 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.
- [26] E. Price, Y. T. Liu, M. J. Black, and A. Ahmad, "Simulation and control of deformable autonomous airships in turbulent wind," in *16th International Conference on Intelligent Autonomous System (IAS)*, June 2021.
- [27] A. Rottmann and W. Burgard, "Adaptive autonomous control using online value iteration with gaussian processes," in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 2106–2111.
- [28] M. Jaderberg, V. Mnih, W. M. Czarnecki, T. Schaul, J. Z. Leibo, D. Silver, and K. Kavukcuoglu, "Reinforcement learning with unsupervised auxiliary tasks," 2016.
- [29] M. Riedmiller, R. Hafner, T. Lampe, M. Neunert, J. Degraeve, T. Wiele, V. Mnih, N. Heess, and J. T. Springenberg, "Learning by playing solving sparse reward tasks from scratch," in *International conference on machine learning*. PMLR, 2018, pp. 4344–4353.
- [30] T. Yu, S. Kumar, A. Gupta, S. Levine, K. Hausman, and C. Finn, "Gradient surgery for multi-task learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5824–5836, 2020.
- [31] Y. Pan, K. Banman, and M. White, "Fuzzy tiling activations: A simple approach to learning sparse representations online," *arXiv preprint arXiv:1911.08068*, 2019.
- [32] S. Sinha, H. Bharadhwaj, A. Srinivas, and A. Garg, "D2rl: Deep dense architectures in reinforcement learning," *arXiv preprint arXiv:2010.09163*, 2020.
- [33] 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.