

Multi-Task Reinforcement Learning in Continuous Control with Successor Feature-Based Concurrent Composition

Yu Tang Liu^{1,2} and Aamir Ahmad^{2,1}

Abstract—Deep reinforcement learning (DRL) frameworks are increasingly used to solve high-dimensional continuous-control tasks in robotics. However, due to the lack of sample efficiency, applying DRL for online learning is still practically infeasible in the robotics domain. One reason is that DRL agents do not leverage the solution of previous tasks for new tasks. Recent work on multi-task DRL agents based on successor features (SFs) has proven to be quite promising in increasing sample efficiency. In this work, we present a new approach that unifies two prior multi-task RL frameworks, SF-GPI and value composition, and adapts them to the continuous control domain. We exploit compositional properties of successor features to compose a policy distribution from a set of primitives without training any new policy. Lastly, to demonstrate the multi-tasking mechanism, we present our proof-of-concept benchmark environments, Pointmass and Pointer, based on IsaacGym, which facilitates large-scale parallelization to accelerate the experiments. Our experimental results show that our multi-task agent has single-task performance on par with soft actor-critic (SAC), and the agent can successfully transfer to new unseen tasks. We provide our code as open-source at https://github.com/robot-perception-group/concurrent_composition for the benefit of the community.

I. INTRODUCTION

One approach to achieving optimal continuous control in robotics is through reinforcement learning-based methods (RL) [1]. However, if the reward definition changes, RL agents are usually optimized for one specific task and require more training. On the other hand, multi-task RL frameworks aim to design an RL agent that allows recycling old policies for other tasks to achieve higher sample efficiency [2].

One way to achieve multi-task RL is through transfer learning by manipulating the policies trained on the old tasks. This work focuses on the transfer learning methods that define tasks using linearly decomposable reward functions, commonly used to define many robot control tasks [3].

Currently, there are two major transfer learning frameworks following the aforementioned assumption. The first framework, using *successor features* and *generalized policy improvement* (SF-GPI, [4], [5]), derives the new policy by choosing the sub-policies, or primitives, with the largest associated action-value. The other framework directly composes the action-value of different tasks, which we refer to as *value composition* (VC [6], [7], [8], [9]). This framework first defines the composite value functions of new tasks by linearly combining the constituent action-value functions and then extracts the policy.

¹Max Planck Institute for Intelligent Systems, 72076 Tübingen, Germany. ²University of Stuttgart, 70569 Stuttgart, Germany. yutang.liu@tuebingen.mpg.de, aamir.ahmad@ifr.uni-stuttgart.de




Fig. 1: In concurrent composition, policy extraction is intractable online. Instead, we propose composing the primitives directly in the run time.

Different from the *Option* framework [10], which is another approach to multi-task RL that combines the primitives in temporal order, in our novel approach, we focus on concurrent policy composition by composing the primitives at every time step. The main objective is to train primitives and enable the online composition to transfer knowledge for solving arbitrary tasks, thereby enhancing sample efficiency [4]. Policy extraction is prohibitive since it requires updating policies multiple times. We consider real-time performance critical. Therefore, we must adapt VC-based methods for online composition. We do so by multiplicative policy composition (MCP [11]). We show this can be achieved via the successor features framework, where SF-GPI becomes one special case.

In summary, our contributions include the following:

- A novel successor feature-based online concurrent composition framework unifying the GPI and VC methods.
- Deriving the relation between composition in value space and policy space.
- Extending this framework to composition in action space via the impact matrix.
- Introducing two benchmark multi-task environments, Pointmass and Pointer, based on IsaacGym [12].

II. RELATED WORK

The compositionality in continuous control was first exploited by linearly-solvable Markov Decision Processes (LMDPs) [13]. It can be shown that composing the optimal value function of LMDPs results in the optimal value function in the composite task [14]. However, LMDPs were restricted to known dynamics.

The compositional agents in our framework instead build upon the two transfer learning frameworks: VC and SF-GPI. The VC-based methods often exploit the property that the optimal value function of conjunction tasks can be well approximated by the average of the constituent value functions [6], [7]. Interestingly, boolean operators can further expand the valid compositions in the discrete setting and achieve a

super-exponential amount of task transfer [8], [9]. However, none of these approaches consider concurrent composition scenarios.

On the other hand, **SF-GPI** [4] applies successor features [15] to estimate a set of Q-functions in the new task. Maximizing these Q-functions results in policy improvement. The results in [4] indicate that simultaneously training several primitives can significantly improve the sample efficiency for solving a new task with collective knowledge. Our work can be seen as an extension to SF-GPI by bringing valid operations from VC to successor feature-enabled concurrent learning.

III. BACKGROUND

A. Multi-Task Reinforcement Learning

RL improves a policy through interaction with the environment, and it formulates this procedure by *Markov decision processes* (MDPs), $\mathcal{M} \equiv \langle \mathcal{S}, \mathcal{A}, p, r, \gamma \rangle$. \mathcal{S} and \mathcal{A} define the state space and action space respectively. The dynamics $p(s_{t+1}|s_t, a_t)$ defines the transition probability from one state to another. Reward function $r(s_t, a_t, s_{t+1})$ governs the desired agent behavior. The discount $\gamma \in [0, 1)$ determines the valid time horizon of the task. The goal is to find the optimal control policy $\pi(a|s) : \mathcal{S} \rightarrow \mathcal{A}$, such that the expected discounted return, $J(\pi) = \mathbb{E}_{\pi, \mathcal{M}}[\sum_{\tau=t}^{\infty} \gamma^{\tau-t} r_{\tau}]$, can be maximized.

We consider transfer within a set of MDPs \mathcal{M}_w where $\mathcal{M}_w \equiv \langle \mathcal{S}, \mathcal{A}, p, r_w, \gamma \rangle$, and each has its unique reward function, $r_w(s_t, a_t, s_{t+1}) = \phi(s_t, a_t, s_{t+1})^\top \cdot w$, which is a linear combination of a set of commonly shared d features, $\phi(s_t, a_t, s_{t+1}) \in \mathbb{R}^d$ and the task weight $w \in \mathbb{R}^d$. The goal is to find a set of primitives that can solve the subset of tasks $w' \in \mathcal{W}'$ such that we can recover the policies of the unseen tasks $w \in \mathcal{W}$ by the found primitives, where $\mathcal{W}' \subset \mathcal{W}$.

B. Maximum Entropy Reinforcement Learning

The maximum entropy RL framework [16] introduces a policy entropy bonus in the reward function. The goal is to find an optimal control policy that not only maximizes the trajectory reward but also encourages policy randomness, i.e.,

$$\pi^* = \operatorname{argmax}_{\pi} \mathbb{E}_{\pi, \mathcal{M}_w} \left[\sum_{\tau=t}^{\infty} \gamma^{\tau-t} (r_{w, \tau} + \alpha \mathcal{H}(\pi(\cdot|s_{\tau}))) \right], \quad (1)$$

where α is a temperature scaler controlling the entropy bonus and $\mathcal{H}(\cdot)$ is an entropy estimator. Prior works have found that the entropy bonus can encourage exploration [17], increase the learning stability in noisy environments [18], and preserve the policy multi-modality [19].

We define the surrogate objective: action-value function, or Q-function, of policy π on task w as:

$$Q_w^\pi(s_t, a_t) \equiv r_{w, t} + \mathbb{E}_{\pi} \left[\sum_{\tau=t+1}^{\infty} \gamma^{\tau-t} (r_{w, \tau} + \alpha \mathcal{H}(\pi(\cdot|s_{\tau}))) \right], \quad (2)$$

and optimizing Q-function will result in an optimal maximum entropy policy. Note that when $\alpha = 0$, we recover the standard RL Q-function. The Q-function can represent the performance of its correspondent policy, which can be improved recursively through the *soft policy iteration*.

C. Soft Policy Iteration

The soft policy iteration can be defined in two steps: (1) *soft policy evaluation* (2) *soft policy improvement* [19]. The evaluation step estimates Q-function of the training policy through temporal difference learning, and it is formulated as follows,

$$Q_w^\pi(s_t, a_t) = r_w(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim p} [V_w^\pi(s_{t+1})] \quad (3)$$

$$V_w^\pi(s_t) = \alpha \log \int_{a \in \mathcal{A}} e^{Q_w^\pi(s_t, a)/\alpha}. \quad (4)$$

Then the desired Boltzmann policy can be updated according to the soft policy improvement step,

$$\pi_{desired}(a_t|s_t) = \frac{e^{Q_w^\pi(s_t, a_t)}}{\int_{a \in \mathcal{A}} e^{Q_w^\pi(s_t, a)}} = \frac{e^{Q_w^\pi(s_t, a_t)}}{Z^\pi(s_t)} \propto e^{Q_w^\pi(s_t, a_t)}, \quad (5)$$

by solving the equation w.r.t each state,

$$\pi(\cdot|s_t) = \operatorname{argmin}_{\pi'} D \left(\pi'(\cdot|s_t) \parallel e^{Q_w^\pi(s_t, \cdot)} \right), \quad (6)$$

where D denotes the Kullback-Leiber divergence. Intuitively, the goal is to match the policy distribution to the Q-value of different actions; therefore, the higher the Q-value, the higher the probability of selecting the action. In practice, the integration above is replaced with the summation for the tractability. And we approximate the Boltzmann policy by other distributions, e.g. a Gaussian $\pi = \mathcal{N}^{|\mathcal{A}|}$.

It can be shown that the evaluation step is a contraction with optimal Q-function as its fixed point, while the policy improvement step results in an improvement policy. Therefore, soft policy iteration will eventually converge to the optimal Q-function Q_w^* on task w with its corresponding optimal policy π^* , where $Q_w^* \equiv Q_w^{\pi^*}$ [19].

D. Successor Feature (SF)

Given that $r_w(s_t, a_t, s_{t+1}) = \phi(s_t, a_t, s_{t+1})^\top \cdot w$, we can rewrite the standard Q-function in following form,

$$Q_w^\pi(s_t, a_t) = \mathbb{E}_{\pi} \left[\sum_{\tau=t}^{\infty} \gamma^{\tau-t} r_{w, \tau} \right] \quad (7)$$

$$= \mathbb{E}_{\pi} \left[\sum_{\tau=t}^{\infty} \gamma^{\tau-t} \phi_t^\top \right] \cdot w = \psi^\pi(s_t, a_t)^\top \cdot w \quad (8)$$

where $\psi^\pi \in \mathbb{R}^d$ is the *successor feature* of policy π [15]. Similar to Q-functions, which approximate "reward-to-go," SFs approximate the "feature-to-go" for a policy. Since Q-function can be viewed as a performance metric of a policy, SFs provide a shortcut to estimate policy performance in different tasks w . In analogous to (3, 4), we can approximate SFs via soft policy iteration [7], i.e.

$$\psi^\pi(s_t, a_t) = \phi_t(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim p} [\Upsilon(s_{t+1})], \quad (9)$$

$$\Upsilon^\pi(s_t) = \alpha \log \int_{\mathcal{A}} e^{\psi^\pi(s_t, a_t)/\alpha}. \quad (10)$$

The policy improvement step is similar to (6) by replacing the Q-function with the approximation from the successor features and the desired primitive tasks, which is defined in Sec. IV-B.

E. Concurrent Composition

During the training process for a specific task w , a single task RL agent always possesses a *true* Q-function denoted as Q_w^{True} . However, Q_w^{True} does not exist in a multi-task setting.

Instead, it needs to be approximated using a set of primitive Q-functions $Q_w^\Pi = \{Q_w^{\pi_1}, \dots, Q_w^{\pi_n}\}$ and value composition rules, i.e., $\hat{Q}_w^{Rule} = Rule(Q_w^\Pi)$, resulting in a composition loss, $|Q_w^{True} - \hat{Q}_w^{Rule}|$. One of our objectives is to identify a rule that minimizes this loss.

Similarly, in continuous control, a single task agent also possesses a *true* policy denoted as π_w^{True} . The objective of policy composition is to approximate the true task policy π_w^{True} by the composite $\hat{\pi}_w^{Rule} = Rule(\Pi, w)$ using a set of primitives $\Pi = \{\pi_1, \dots, \pi_n\}$ and composition rules. We will show that for certain compositions in value space, corresponding rules exist in policy space. Furthermore, the composition loss in value space can lead to a performance loss in the composite policy, i.e., $Q_w^{\hat{\pi}_w^{Rule}} \leq Q_w^{\pi_w^{True}}$. This insight is crucial: the composite policy can only achieve performance equal to or lower than the single-task policy due to the incurred composition loss.

F. Generalized Policy Improvement Composition (GPI)

When provided with primitives Π and their corresponding Q-functions in a task w , denoted as Q_w^Π , a composite value, and policy can be constructed as follows:

$$Q_w^{True}(s, a) \approx \hat{Q}_w^{GPI}(s, a) = \max_{\pi \in \Pi} Q_w^\pi(s, a), \quad (11)$$

$$\hat{\pi}^{GPI}(s) \in \operatorname{argmax}_{a \in \mathcal{A}} \hat{Q}_w^{GPI}(s, a). \quad (12)$$

According to the GPI theorem [4], this new policy $\hat{\pi}^{GPI}$ is guaranteed to be at least as good as the primitives, as its Q-function is always greater than or equal to that of its members: $Q_w^{\hat{\pi}^{GPI}}(s, a) \geq \sup_{\pi \in \Pi} Q_w^\pi(s, a)$.

Successor features offer a quick and convenient method for evaluating the performance of a policy π in any task w , i.e., $Q_w^\pi = w \cdot \psi^\pi$. This evaluation step is known as *generalized policy evaluation* (GPE), and GPI can subsequently be applied to select the action that yields the highest value.

Note that our SF-GPI (Sec. IV-C.3) sample actions from the primitives, i.e. $\operatorname{argmax}_{a \sim \Pi}$, since the $\operatorname{argmax}_{a \in \mathcal{A}}$ expression is infeasible in continuous action spaces.

G. Value Composition (VC)

Combining the Q-functions provides a means to obtain the Q-function for a new task without explicitly solving for it. Previous studies have established a set of valid operations for multi-goal discrete control [8], [9]. In continuous control, our focus is linear task combinations that do not assume binary task weights. Specifically, if a new task is a linear combination of old tasks, denoted as $r_{new} = b \cdot r_0 + (1 - b) \cdot r_1$, the optimal constituent Q-functions can approximate the optimal Q-function of the new task [6]: $Q_{new}^* \approx \hat{Q}_{new}^\Sigma = b \cdot Q_0^* + (1 - b) \cdot Q_1^*$. The new policy can then be extracted from the composite. However, in continuous control, extracting process requires several policy improvement steps, making it intractable online. To extend this framework to concurrent composition, we relax the assumption and generalize the theorem to non-optimal value functions: $Q_{new}^{True} \approx \hat{Q}_{new}^\Sigma = b \cdot Q_0^{\pi_0} + (1 - b) \cdot Q_1^{\pi_1}$. With 8, we can formulate VC to SFs framework.

H. Multiplicative Compositional Policy (MCP)

MCP [11] has shown to compose the policy primitive in the imitation learning scenario. Here, we introduce it to concurrent composition. Given a set of k Gaussian primitives $\Pi \in \mathcal{N}^{|\mathcal{A}| \times k}$ and a gating vector $g \in [0, 1]^k$ where $\|g\|_1 = 1$, the function $MCP(\Pi, g) : \mathcal{N}^{|\mathcal{A}| \times k} \times \mathbb{R}^k \rightarrow \mathcal{N}^{|\mathcal{A}| \times 1}$ is defined as follows,

$$\pi(a|s) = MCP(\Pi, g) \equiv \frac{1}{Z(s)} \prod_{i=1}^k \pi_i(a|s)^{g_i}, \quad (13)$$

where $Z(s)$ is a partition function for normalization. If all the primitives are Gaussian, then the composite is also a Gaussian. It was not explained why such composition is valid. However, we show this is a natural outcome from value composition in the policy space (Sec. IV-C.1).

IV. METHODOLOGY

We present our novel approach in the following sections, which are eventually summarized in Algorithm 5.

A. Successor Feature-based Composition

We introduce two more SF-based compositions and our GPI composition rule (Sec. IV-C.3).

1) *Successor Feature based Value Composition (SFV)*: First, we introduce the linear value composition to the successor feature framework. Given the task $r_w = \phi^\top \cdot w$, we can treat features ϕ as sub-tasks. From this perspective, ϕ becomes the constituent sub-tasks, and r_w is the composite task. Thus, the successor features represent the primitive's performance on each sub-task, e.g., $\psi_j^{\pi_i}$ is the performance metric of π_i on controlling sub-task ϕ_j . Therefore, (8) can be seen as a special case for the value composition of the same policy.

Value composition theorems state that the average of the Q-functions is a sufficiently good approximation of the true Q-function [6], [8], [9]. We can define the composite Q-function and discrete greedy policy accordingly,

$$Q_w^{True}(s, a) \approx \hat{Q}_w^{SFV} \equiv \frac{1}{n} \sum_{i=1}^n Q_w^{\pi_i}(s, a) = \frac{w^\top}{n} \sum_{i=1}^n \psi^{\pi_i}(s, a) \quad (14)$$

$$\hat{\pi}^{SFV}(s) \in \operatorname{argmax}_{a \in \mathcal{A}} \hat{Q}_w^{SFV}(s, a). \quad (15)$$

This expression shows that SFV composition aims to achieve an average performance in all sub-tasks.

2) *Maximum Successor Feature Composition (MSF)*: We can treat successor features as sub-tasks. In that case, it naturally raises the question: can we combine primitives in a certain way to achieve the current best achievable performance in all sub-tasks and, subsequently, any composite task? That is, given SFs of their corresponding n primitives, $\Psi^\Pi = \{\psi^{\pi_1}, \dots, \psi^{\pi_n}\}$, where $\psi^{\pi_i} = \{\psi_1^{\pi_i}, \dots, \psi_d^{\pi_i}\}$. Then the desired composition operator \mathcal{U} can be defined as $\mathcal{U}\Psi^\Pi \equiv \{\max_{\pi \in \Pi} \psi_1^\pi, \dots, \max_{\pi \in \Pi} \psi_d^\pi\}$. With this, we define the composition and its discrete greedy policy

$$Q_w^{True}(s, a) \approx \hat{Q}_w^{MSF}(s, a) \equiv w^\top \mathcal{U}\Psi^\Pi(s, a) \quad (16)$$

$$\hat{\pi}^{MSF}(s) \in \operatorname{argmax}_{a \in \mathcal{A}} \hat{Q}_w^{MSF}(s, a). \quad (17)$$

To summarize, all the proposed composition rules either overestimate or underestimate the true Q-function. Specifically,

they exhibit the following relationships:

$$\hat{Q}_w^{MSF} \geq \hat{Q}_w^{SFV} \geq Q_w^{True} \geq \hat{Q}_w^{GPI}. \quad (18)$$

Like VC, both SFV and MSF assume that all sub-tasks can be accomplished simultaneously, leading to an overestimation when this assumption does not hold. Conversely, GPI assumes that only one task can be achieved at a time, resulting in an underestimation. Due to the imperfect nature of composition rules, the composites are rarely optimal, i.e.,

$$Q_w^{True} \geq \{Q_w^{SFV}, \hat{Q}_w^{MSF}, Q_w^{GPI}\}. \quad (19)$$

B. Construct Primitives

We train primitives based on the sub-tasks, and each primitive is only responsible for its sub-task. Therefore, we have the same number of sub-tasks and primitives (i.e., $d = n$). All primitives are trained by soft policy iterations and follow its convergence guarantee, i.e.

$$\pi_i(\cdot|s) \propto e^{Q_{e_i}^{\pi_i}(s,\cdot)} = e^{\mathbf{e}_i^\top \cdot \psi^{\pi_i}(s,\cdot)}. \quad (20)$$

The base tasks $\mathbf{e}_i \in \mathbb{R}^d$ can be constructed arbitrarily, as long as the base task and the transfer task share the same sub-tasks or overlap in value space. For example, the base task defined as the zero vector will fail since this task has no overlapping with any transfer task. For simplicity, we define \mathbf{e}_i by elementary vector such that the primitives form the "basis" in the policy space.

C. Approximate True Policy by Composing Primitives

This is the core part of our method. We provide an analytical view of how the true policy can be reconstructed by a set of primitives given their relation in value space introduced in Sec. IV-A.

1) *MSF Composition*: Following MSF in value space (17), we derive the correspondent policy composition,

$$\pi^{True}(\cdot|s) \propto e^{Q_w^{True}(s,\cdot)} \approx e^{\hat{Q}_w^{MSF}(s,a)} = e^{w^\top \mathcal{U}_{\Psi^\Pi}(s,a)} \quad (21)$$

$$\approx e^{w^\top \cdot [\psi^{\pi_1}, \dots, \psi^{\pi_d}](s,\cdot)} \propto \prod_{i=1}^d \pi_i(\cdot|s)^{w_i}. \quad (22)$$

The approximation in (22) assumes the maximum SF of each sub-task is most likely to be the SF of the primitive trained from the respective sub-task. Then the MSF composition can be achieved by MCP, i.e., $\hat{\pi}(\cdot|s) = MCP(\Pi, g)$, where $g = \frac{w}{|w|}$ is the normalized gating vector.

2) *SFV Composition*: Following SFV in value space (14), we derive the correspondent policy composition,

$$\pi^{True}(\cdot|s) \propto e^{Q_w^{True}(s,\cdot)} \approx e^{\hat{Q}_w^{SFV}(s,\cdot)} = \prod_{i=1}^d e^{\frac{w_i^\top}{d} \psi^{\pi_i}(s,\cdot)} \quad (23)$$

$$= \prod_{i=1}^d e^{\frac{1}{d} w_i \psi_i^{\pi_i}(s,\cdot)} \prod_{i=1}^d \prod_{j \neq i}^d e^{\frac{1}{d} w_j \psi_j^{\pi_i}(s,\cdot)} \propto \prod_{i=1}^d \pi_i(\cdot|s)^{\frac{w_i}{d}} \cdot \epsilon \quad (24)$$

where ϵ is the artefactual noise. Because of the noise, the SFV composition is valid but impractical in policy space with the set of primitives introduced in Sec. IV-B.

3) *GPI Composition*: Though this is already an SF-based composition, we reformulate it for completeness.

$$\pi^{True}(\cdot|s) \propto e^{Q_w^{True}(s,\cdot)} \approx e^{\hat{Q}_w^{GPI}(s,\cdot)} \quad (25)$$

$$= e^{\max_{\pi_i \in \Pi} Q_w^{\pi_i}(s,\cdot)} \propto \operatorname{argmax}_{\pi_i \in \Pi} Q_w^{\pi_i}(s,\cdot) \quad (26)$$

The approximation can be achieved by successor features, i.e. $\hat{\pi}(\cdot|s) = \operatorname{argmax}_{\pi_i \in \Pi} Q_w^{\pi_i}(s,\cdot) = \operatorname{argmax}_{\pi_i \in \Pi} w^\top \psi^{\pi_i}(s,\cdot)$.

We summarize the above policy composition rules in Algorithms 1 2. For clarity, our proposed approach incorporates two transfer mechanisms that complement each other to expedite learning. First, SFs facilitate transfer to any task by specifying different task weights, enabling policy evaluation without additional training. This transfer is associated with composition in the value space. Second, as the composite policy outperforms the primitives, the trajectories gathered by the composite explore higher-value regions in the composite task. This transfer mechanism resembles *teacher-student* [20] or *learn-from-demonstration* [21] approaches, where knowledge is transferred from the expert policy to the primitives.

At this point, we have constructed an analytic method to find the policy composition corresponding to the composition in value space. Note that if primitives are constructed with different base tasks \mathbf{e}_i , the composition in policy space will change accordingly. Next, we heuristically extend the idea to *direct action composition*, which allows composition to happen in the action space.

Algorithm 1 MSF Composition

- 1: **function** MSF(s, w):
 - 2: $\hat{\pi}(\cdot|s) = MCP(\Pi(\cdot|s), \frac{w}{|w|})$
 - 3: $a \sim \hat{\pi}(\cdot|s)$ ▷ sample from composite
 - 4: **return** a
-

Algorithm 2 GPI Composition

- 1: **function** SF-GPI(s, w):
 - 2: $\mathbf{a} \sim \Pi(\cdot|s)$ ▷ sample from primitives
 - 3: $\hat{Q}_w^\pi = \psi^\pi(s, \mathbf{a}) \cdot w$ ▷ GPE
 - 4: $a = \operatorname{argmax}_{a \sim \Pi} \hat{Q}_w^\pi$ ▷ GPI
 - 5: **return** a
-

D. Multiplicative Compositional Action (MCA)

Rewriting a Gaussian primitive $\pi_i \in \Pi$ in the form,

$$\pi_i(a|s) = [\mathcal{N}_1^{\pi_i}, \dots, \mathcal{N}_{|\mathcal{A}|}^{\pi_i}] \quad (27)$$

where $\mathcal{N}_j^{\pi_i} = \mathcal{N}^{\pi_i}(\mu_j, \sigma_j)$ is the shorthand for the j-th action component of π_i . Then the MSF composition (22) has the action components in the following form,

$$\hat{\pi}(\cdot|s) = MCP(\Pi, g) = \left[\prod_{i=1}^d (\mathcal{N}_1^{\pi_i})^{g_i}, \dots, \prod_{i=1}^d (\mathcal{N}_{|\mathcal{A}|}^{\pi_i})^{g_i} \right]. \quad (28)$$

However, since all action components are coupled by the exponent g_i , we cannot manipulate them differently. Alternatively, given $\kappa \in [0, 1]^{|\mathcal{A}| \times n}$ where $\|\kappa\|_1 = 1$, we define MCA that applies composition in the action space

$MCA(\Pi, \kappa) : \mathcal{N}^{|\mathcal{A}| \times n} \times \mathbb{R}^{|\mathcal{A}| \times n} \rightarrow \mathcal{N}^{|\mathcal{A}| \times 1}$;

$$\hat{\pi}(\cdot|s) = MCA(\Pi, \kappa) \equiv \left[\prod_{i=1}^d (\mathcal{N}_1^{\pi_i})^{\kappa_1^{\pi_i}}, \dots, \prod_{i=1}^d (\mathcal{N}_{|\mathcal{A}|}^{\pi_i})^{\kappa_{|\mathcal{A}|}^{\pi_i}} \right]. \quad (29)$$

We define $\kappa_j^{\pi_i}$ in such a way that action components with lower successor feature value will be filtered out and higher value to be emphasized, i.e., $\kappa^{\pi_i} \propto \mathcal{P}_{|\mathcal{A}| \times d} \cdot \psi^{\pi_i}$, where \mathcal{P} is an *impact matrix* defining the mapping from successor feature to the action components.

E. Impact Matrix

We would like to know the relationship between each action component and the features throughout the trajectory, which is, surprisingly, the successor feature. One way to determine the relationship between two variables is by observing the change in one variable when the other is tuned, which can be achieved through derivatives. Given a set of successor features $\psi = \{\psi_1, \dots, \psi_d\}$, we define the impact matrix as the absolute value of the Jacobian of ψ w.r.t the action components:

$$\mathcal{P} \equiv |\nabla_a \psi(s, a)| = \begin{bmatrix} \left| \frac{\partial \psi_1}{\partial a_1} \right| & \dots & \left| \frac{\partial \psi_d}{\partial a_1} \right| \\ \vdots & \ddots & \vdots \\ \left| \frac{\partial \psi_1}{\partial a_{|\mathcal{A}|}} \right| & \dots & \left| \frac{\partial \psi_d}{\partial a_{|\mathcal{A}|}} \right| \end{bmatrix}_{|\mathcal{A}| \times d} \quad (30)$$

The computation required by this expression and the accuracy of the derivative are major concerns. Fortunately, in practice, this estimation can be efficiently computed using vectorized operations with GPUs. The accuracy of the derivative improves over time as the successor features become smoother during training. However, in the early stages of training, the estimation can be noisy. We can mitigate this noise by averaging over all policies, i.e., $\psi \approx \frac{1}{n} \sum_n \psi^{\pi_n}$. Consequently, the estimation becomes less noisy with an increase in the number of primitives. To further reduce noise, we incorporate temporal dependency by averaging with the impact matrix from the previous time step, i.e., $\mathcal{P} \approx (\mathcal{P}_{t-1} + \mathcal{P}_t)/2$.

F. Direct Action Composition (DAC)

With the impact matrix, we define $\kappa \equiv \mathcal{P}_{|\mathcal{A}| \times d} \cdot \Psi_{d \times n} \odot w_{d \times 1}$, where

$$\Psi = \{\psi^{\pi_1}, \dots, \psi^{\pi_n}\} = \begin{bmatrix} \psi_1^{\pi_1} & \dots & \psi_1^{\pi_n} \\ \vdots & \ddots & \vdots \\ \psi_d^{\pi_1} & \dots & \psi_d^{\pi_n} \end{bmatrix}_{d \times n}, \quad (31)$$

\odot denotes the Hadamard product, and $n = d$. Intuitively, task weights w filter the task-relevant successor features and then map their impact to each action component via the impact matrix, i.e., $\hat{\kappa} = \mathcal{P} \cdot \Psi \odot w$. And then DAC can be achieved by MCA (29) as shown in Algorithms 3. We introduce another variant, DAC-GPI, by choosing the action component with the largest $\hat{\kappa}$, as displayed in Algorithms 4.

In practice, we replace Ψ with advantage Γ for numerical stability, i.e. $\Gamma = \Psi - \frac{1}{n} \sum_n \psi^{\pi_n}$. Empirically, we found the performance much better when we clip the advantage below zero, i.e., $\Gamma_{clip} = \max(0, \Gamma)$. Finally, we have $\hat{\kappa} = \mathcal{P} \cdot \Gamma_{clip} \odot w$.

Notice that DAC-GPI reduces to SF-GPI when all the action components come from the same policy. Moreover, when

Algorithm 3 DAC Composition

```

1: function DAC(s, w):
2:    $\mathbf{a} \sim \Pi(\cdot|s)$  ▷ sample from primitives
3:    $\hat{\kappa} = \mathcal{P} \cdot \Psi(s, \mathbf{a}) \odot w$ 
4:    $\hat{\pi}(\cdot|s) = MCA(\Pi(\cdot|s), \frac{\hat{\kappa}}{|\hat{\kappa}|})$ 
5:    $a \sim \hat{\pi}(\cdot|s)$  ▷ sample from composite
6:   return  $a$ 

```

Algorithm 4 DAC-GPI Composition

```

1: function DAC-GPI(s, w):
2:    $\mathbf{a} \sim \Pi(\cdot|s)$  ▷ sample from primitives
3:    $\hat{\kappa} = \mathcal{P} \cdot \Psi(s, \mathbf{a}) \odot w$ 
4:    $a = [\operatorname{argmax}_{a_1} \hat{\kappa}_1^\pi, \dots, \operatorname{argmax}_{a_{|\mathcal{A}|}} \hat{\kappa}_{|\mathcal{A}|}^\pi]$  ▷ GPI
5:   return  $a$ 

```

each primitive shares the same exponent across all action components, DAC reduces to MSF. The major difference between DAC and DAC-GPI is that the former composes action components from all primitives according to the SFs, while the latter only selects the best.

In summary, we have constructed a method to bridge the composition in value space to policy space. This provides an effective analytic method for generating different concurrent composition algorithms. We then show that the composition can happen in policy space and action space with the help of the impact matrix. Unlike standard RL agents, compositional agents usually involve critics or SFs in the action selection process and filter out the lower-quality actions before delivery.

G. Training the Network

The training data is gathered by direct interaction with the environment through composition rules. Therefore, we can treat the concurrent compositional agent just as any regular RL agent and train end-to-end.

In total, we have n pairs of SFs and primitives with the i -th successor feature network ψ^{π_i} , target network $\bar{\psi}^{\pi_i}$, and the primitive network π_i parameterized by $\theta_{\psi^{\pi_i}}$, $\theta_{\bar{\psi}^{\pi_i}}$, and θ_{π_i} respectively. The training scheme follows the soft policy iteration, which includes the evaluation step (9,10) and improvement step (6). The evaluation step now evaluates the successor feature instead of the Q-function. In contrast, the policy improvement step stays the same by following the base task Q-function $Q_{\theta_i}^{\pi_i}$. Therefore, the loss for the i -th successor feature can be computed as follows,

$$J_{\psi^{\pi_i}}(\theta_{\psi^{\pi_i}}) = \mathbb{E}_{(s_t, a_t) \sim \mathcal{D}} \left[(\psi^{\pi_i}(s_t, a_t) - \hat{\psi}^{\pi_i}(s_t, a_t))^2 \right], \quad (32)$$

where \mathcal{D} is the replay buffer and we have the TD-target $\hat{\psi}^{\pi_i}$ approximated by the target network $\bar{\psi}^{\pi_i}$,

$$\hat{\psi}^{\pi_i}(s_t, a_t) = \phi_t + \gamma \mathbb{E}_{s_{t+1} \sim p(\cdot|s_t, a_t)} [\mathbf{Y}^{\pi_i}(s_{t+1})], \quad (33)$$

$$\mathbf{Y}^{\pi_i}(s_t) = \mathbb{E}_{a_t \sim \pi_i} [\bar{\psi}^{\pi_i}(s_t, a_t) - \log \pi_i(a_t|s_t)]. \quad (34)$$

where $-\alpha \log \pi$ is a one-step policy entropy estimator. Then the primitive can be improved by reducing the KL-divergence in (6), which can be approximated by maximizing the

objective,

$$J_{\pi_i}(\theta_{\pi_i}) = \mathbb{E}_{s_t \sim \mathcal{D}, a_t \sim \pi_i} [Q_{e_i}^{\pi_i}(s_t, a_t) - \log \pi_i(a_t | s_t)], \quad (35)$$

with the primitive Q-function computed by (20). We summarize our approach in Algorithm 5. The method alternates between the interaction phase for collecting samples by the compositional policy and the learning phase that updates the function approximators' parameters by stochastic gradient descent using the collected samples. Although looping over all primitives may appear intimidating in the algorithm, in our implementation, we vectorized computations to update all primitives simultaneously in a single pass such that the computation time remains constant.

Algorithm 5 Composition Agent

Choose *method* \in MSF, SF-GPI, DAC, DAC-GPI
Initialize network parameters $\theta_{\psi^\pi}, \theta_{\bar{\psi}^\pi}, \theta_\pi$

- 1: **for** N steps **do**
- 2: **if** exploration **then** $a \leftarrow \text{Uniform}(\mathcal{A})$
- 3: **else** $a \leftarrow \text{method}(s)$
- 4: $\mathcal{D} \leftarrow \mathcal{D} \cup (s_t, a_t, \phi(s_t, a_t), s_{t+1})$
- 5: **for** each gradient step **do**
- 6: **for** $i \leftarrow 1, 2, \dots, n$ **do**
- 7: $\theta_{\psi^{\pi_i}} \leftarrow \theta_{\psi^{\pi_i}} - \lambda_\psi \nabla_{\theta_{\psi^{\pi_i}}} J_{\psi^{\pi_i}}(\theta_{\psi^{\pi_i}})$
- 8: $\theta_{\pi_i} \leftarrow \theta_{\pi_i} - \lambda_\pi \nabla_{\theta_{\pi_i}} J_{\pi_i}(\theta_{\pi_i})$
- 9: $\theta_{\bar{\psi}^{\pi_i}} \leftarrow \tau \theta_{\bar{\psi}^{\pi_i}} + (1 - \tau) \theta_{\psi^{\pi_i}}$

V. EXPERIMENTS AND RESULTS

The experiment aims to test if the multi-task RL agents generated by the proposed framework, i.e., SF-GPI, MSF, DAC, DAC-GPI, can be task-agnostic and transfer to unseen tasks.

A. Network Architecture

We represent primitives and successor features by two neural networks as displayed in Fig. 2. Knowledge transfer can happen within the network since the layers or representations are shared. To reduce over-estimation in the SFs and enable




Fig. 2: network architecture.

higher update-to-data ratio (UTD ratio) [22], we include *dropout* and *layernorm* to the SF network. To further mitigate the positive bias [23], we apply the minimum double critic technique for the SF estimation, i.e., $\hat{\psi}^\pi = \min(\psi_{\theta_1}^\pi, \psi_{\theta_2}^\pi)$.

B. Environment and Tasks

Our training environment is implemented using IsaacGym, leveraging its extensive parallelization capabilities. While IsaacGym offers a range of pre-existing robots, we introduce abstract omni- and uni-directional robots, namely Pointmass and Pointer, designed to tackle our problem and gain valuable insights.

1) *State, Action, and Feature Space*: The state and action space of the environments are presented in Table. 1,

	Pointmass	Pointer
\mathcal{A}	(a_x, a_y, a_z)	$(a_{thrust}, a_{roll}, a_{pitch}, a_{yaw})$
\mathcal{S}	$\Delta\mathbb{X}, \Delta\mathbb{V}, \Delta\Theta, \Delta\Omega \in \mathbb{R}^3$	

TABLE 1: The action space for Pointmass are applied force on the x-, y-, and z-axis while for Pointer are applied thrust and attitude controls. Both environments share the same state space, namely relative position, velocity, orientation, and angular velocity between the robot and goal.

Features, or sub-tasks, can be defined arbitrarily as long as they can be derived from the state space. We define three sets of features for the Pointmass and one for the Pointer, as shown in Table. 2. The feature sets used in our study include the following variations: *Regular*, which represents a standard approach for defining reward functions in robotics; *Simple*, a simplified version used for demonstration purposes; and *Augment*, which includes redundant features that generate additional primitives. For example, both $-|\Delta\mathbb{X}|$ and $-||\Delta\mathbb{X}||_2$ can achieve position control. Each primitive aims to maximize its corresponding feature or sub-task. In general, the task can be solved only when a sufficient amount of features are provided.


Pointmass	Regular	$- \Delta\mathbb{X} , - \Delta\mathbb{V} _2, Success$
	Simple	$- \Delta\mathbb{X} $
	Augment	$- \Delta\mathbb{X} , - \Delta\mathbb{X} _2, - \Delta\mathbb{V} _2, - \Delta\mathbb{V} _2, Success$
Pointer	Regular	$- \Delta\mathbb{X} _2, - \Delta\mathbb{V} _2, - \Delta\Theta _2, - \Delta\Omega _2, Success$

TABLE 2: Environment feature sets. The absolute value and L2 norm are denoted as $|\cdot|, ||\cdot||_2$, respectively. *Success* is a binary feature triggered when the relative distance is less than 1 meter, i.e., $Success = 1$ if $||\mathbb{X}||_2 < 1$ else 0. The default will be *Regular* if not explicitly mentioned.


2) *Mixture Task*: We introduce a mixture task, which includes two tasks: *navigation* and *hover* tasks. The navigation task aims to bring the robot closer to the goal position, while the hover task requires the robot's velocity to approach the target velocity when near the goal. The task switches from navigation to hover when the distance to the goal is less than three meters. We can formulate this accordingly,

$$w^{mix} = \begin{cases} w^{nav} & \text{if } d(s_{robot}, s_{target}) \geq 3m \\ w^{hov} & \text{otherwise.} \end{cases} \quad (36)$$


where by default, in Pointmass, w^{mix} is set as $[(1, 0, 10), (0, 1, 10)]$ (bold symbols are vectors of the same length to the corresponding features). In Pointer, the default w^{mix} is $[(1, .5, 0, 0, 10), (0, 1, 0, 0, 10)]$. The




(a) **Transfer performance:** Train with a velocity control task and evaluate with a mixture task w^{mix} (36).




(b) **Effect of feature sets in transfer:** Compared to regular features (solid), using augmented features (dotted) is beneficial.




(c) **Single task performance:** Due to the composition loss, composition agents have lower performance compared to the SAC in a single task $w = (1, 0, 1)$.



(d) **Transfer performance:** Train with a velocity control task and evaluate with a mixture task w^{mix} (36).




(e) **Transfer performance:** Train with a velocity control task but evaluate with a position control task $w = (1, 0, 0, 0, 10)$.



(f) **Single task performance:** Due to the composition loss, composition agents have lower performance compared to the SAC in a single task $w = (1, 0, 0, 0, 1)$.

Fig. 3: Empirical results on Pointmass (top row) and Pointer (bottom row) show that the proposed composition agents can gradually improve transfer performance to unseen tasks and have comparable single-task performance to baseline SAC [17]. SAC is presented in transfer tasks to show that tasks are not generalizable. Each curve represents the mean with two standard deviations of 5 experiments run by the best model from the hyper-parameter tuning. The tuning process is conducted by Bayesian optimization with hundred searches for each agent. Our ablation study indicates that dropout can hurt the performance while other techniques' effects remain unclear, including layer norm, prioritized experience replay, entropy tuning, and activation function.



(a) Sample the end position of 20 trajectories for x- (blue), y- (yellow) primitives and their composite (green).

(b) Sample one trajectory for the position (blue), angle (yellow) primitives, and their composite (green).

Fig. 4: Visualize primitive and composite distributions by sample trajectories starting from the initial position X_0 (red) to the goal at the origin. Composition agents allow solving new tasks by composing existing primitives.


navigation task is essentially a position control task. In contrast, in hover tasks, the robots must maintain a circular motion to stay within the hover range and track the desired velocity simultaneously.

We assess transferability by employing different tasks for the training and evaluation phases. Specifically, we train primitives in a velocity control task, with $w_{train} = (0, 1, 0)$ for Pointmass and $w_{train} = (0, 1, .5, 0, 0)$ for Pointer. Then, we evaluate the agents on the mixture task (Fig. 3a 3d) or a position control task (Fig.3e). We observe that multi-task agents demonstrate greater resilience when the evaluation task differs from the training task. The impact of feature sets is illustrated in Fig. 3b, where we find that using the Regular feature set is sufficient to solve the task, and additional features can further enhance transfer performance for MSF and DAC-GPI. Finally, Fig. 4 demonstrates the effectiveness of primitive composition in solving composite tasks that

individual primitives alone cannot handle. We conclude that compositional agents transfer to different tasks if primitives are adequately trained. However, there is no performance guarantee for single task or transfer task due to the composition loss.

composite	action components (a_x, a_y)
$\hat{\pi}^{MSF}$	$(\mathcal{N}_x^{\pi_x} \epsilon, \mathcal{N}_y^{\pi_y} \epsilon)$
$\hat{\pi}^{GPI}$	$(\mathcal{N}_x^{\pi_x} \epsilon, \mathcal{N}_y^{\pi_y} \epsilon)$ or $(\mathcal{N}_x^{\pi_x}, \mathcal{N}_y^{\pi_y} \epsilon)$
$\hat{\pi}^{DAC}$	$((\mathcal{N}_x^{\pi_x})^{\kappa_x} \epsilon) (\mathcal{N}_x^{\pi_x}, (\mathcal{N}_y^{\pi_y})^{\kappa_y} \epsilon) (\mathcal{N}_y^{\pi_y})$

TABLE 3: Composite distributions in Pointmass2D-Simple. ϵ is the composition noise of unknown distribution.



(a) Top row: action components in $\pi_x = (a_x, a_y)$. Bottom row: π_y . Color bar range in $[-1, 1]$.
 (b) Top row: kappa $\kappa^{\pi_x} = (\kappa_x^{\pi_x}, \kappa_y^{\pi_x})$. Bottom row: κ^{π_y} . Color bar range in $[0, 1]$.

Fig. 5: In Pointmass2D-Simple, DAC effectively removes compositional noise (Fig.(a) top right and bottom left) by reducing the corresponding κ . Each sub-figure represents a 2D plane within $[-20, 20]$ meters.

C. Composition Noise and Training Stability in Pointmass2D

The presence of composition noise resulting from redundant action components in the primitive task reduces the training stability. In the Pointmass2D-Simple environment, the primitive actions corresponding to the features $-|\Delta X|$ and $-|\Delta Y|$ generate the distributions $\pi_x = (\mathcal{N}_x^{\pi_x}, \epsilon)$ and $\pi_y = (\epsilon, \mathcal{N}_y^{\pi_y})$, introducing composition noise ϵ (Fig. 5a) and composites (Table. 3). In the case of MSF, the composite action contains noise in both components, while SF-GPI only contains noise in one component. Notably, DAC(-GPI) effectively mitigate noise in both action components, as the successor features reduce the impact of the noisy term (Fig. 5b). The importance of filtering out composition noise for effective transfer is evident in Fig. 3a. DAC(-GPI) exhibits the best asymptotic performance, followed by SF-GPI and MSF. We observed that composition noise hinders the MSF training stability (Fig. 3c, 3f).

We conclude that applying MSF is impractical. Unlike other methods, which can handle any number of primitives, MSF requires that the feature dimension matches the number of primitives (i.e., $n=d$). Moreover, its optimistic assumption of linearity in both task and policy spaces is rarely valid.

D. Single Task Performance

Compared to baseline SAC [17], the multi-task agents exhibit lower asymptotic performance in the single task due to the composition loss (Fig. 3c, 3f). We found that SF-GPI tends to explore less, potentially leading to faster yet premature convergence. We observed that the impact matrix estimation is often noisy, limiting DAC(-GPI)'s effectiveness. Hence, it may be beneficial to initially train primitives using alternative composition methods, transitioning to DAC(-GPI) once SFs achieve satisfactory precision and smoothness.

VI. CONCLUSIONS AND FUTURE WORK

Our work unifies the SF-GPI and value composition to the continuous concurrent composition framework and allows the reconstruction of true policy from a set of primitives. The proposed method was extended to composition in the action space. In the Pointmass and Pointer environments, we demonstrate that our multi-task agents can approximate the true policy in real time and transfer the skills to solve unseen tasks while the single-task performance is competitive with SAC. Our framework incorporates well with the reward-shaping techniques [3] and provides fine control over each sub-tasks via primitives. In addition, the task-agnostic property should benefit the autotelic agent [24] who can set goals and curriculum for themselves [25].

REFERENCES

- [1] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [2] A. Lazaric, "Transfer in reinforcement learning: a framework and a survey," in *Reinforcement Learning: State-of-the-Art*. Springer, 2012, pp. 143–173.
- [3] A. Y. Ng, D. Harada, and S. Russell, "Policy invariance under reward transformations: Theory and application to reward shaping," in *ICML*, vol. 99. Citeseer, 1999, pp. 278–287.

- [4] 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. 30 079–30 087, 2020.
- [5] A. Barreto, D. Borsa, S. Hou, G. Comanici, E. Aygün, P. Hamel, D. Toyama, S. Mourad, D. Silver, D. Precup, *et al.*, "The option keyboard: Combining skills in reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [6] T. Haarnoja, V. Pong, A. Zhou, M. Dalal, P. Abbeel, and S. Levine, "Composable deep reinforcement learning for robotic manipulation," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 6244–6251.
- [7] J. Hunt, A. Barreto, T. Lillicrap, and N. Heess, "Composing entropic policies using divergence correction," in *International Conference on Machine Learning*. PMLR, 2019, pp. 2911–2920.
- [8] B. Van Nieuwerkerk, S. James, A. Earle, and B. Rosman, "Composing value functions in reinforcement learning," in *International conference on machine learning*. PMLR, 2019, pp. 6401–6409.
- [9] G. Nangue Tasse, S. James, and B. Rosman, "A boolean task algebra for reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9497–9507, 2020.
- [10] 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.
- [11] 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.
- [12] J. Liang, V. Makoviychuk, 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.
- [13] E. Todorov, "Linearly-solvable markov decision problems," *Advances in neural information processing systems*, vol. 19, 2006.
- [14] —, "Compositionality of optimal control laws," *Advances in neural information processing systems*, vol. 22, 2009.
- [15] P. Dayan, "Improving generalization for temporal difference learning: The successor representation," *Neural computation*, vol. 5, no. 4, pp. 613–624, 1993.
- [16] B. D. Ziebart, *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. Carnegie Mellon University, 2010.
- [17] 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.
- [18] R. Fox, A. Pakman, and N. Tishby, "Taming the noise in reinforcement learning via soft updates," *Conference on Uncertainty in Artificial Intelligence*. AUAI Press, 2015.
- [19] T. Haarnoja, H. Tang, P. Abbeel, and S. Levine, "Reinforcement learning with deep energy-based policies," in *International conference on machine learning*. PMLR, 2017, pp. 1352–1361.
- [20] S. Schmitt, J. J. Hudson, A. Zidek, S. Osindero, C. Doersch, W. M. Czarnecki, J. Z. Leibo, H. Kuttler, A. Zisserman, K. Simonyan, *et al.*, "Kickstarting deep reinforcement learning," *arXiv preprint arXiv:1803.03835*, 2018.
- [21] T. Hester, M. Vecerik, O. Pietquin, M. Lanctot, T. Schaul, B. Piot, D. Horgan, J. Quan, A. Sendonaris, I. Osband, *et al.*, "Deep q-learning from demonstrations," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [22] T. Hiraoka, T. Imagawa, T. Hashimoto, T. Onishi, and Y. Tsuruoka, "Dropout q-functions for doubly efficient reinforcement learning," *International Conference on Learning Representations (ICLR)*, 2022.
- [23] S. Fujimoto, H. Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *International conference on machine learning*. PMLR, 2018, pp. 1587–1596.
- [24] C. Colas, T. Karch, O. Sigaud, and P.-Y. Oudeyer, "Autotelic agents with intrinsically motivated goal-conditioned reinforcement learning: a short survey," *Journal of Artificial Intelligence Research*, vol. 74, pp. 1159–1199, 2022.
- [25] S. Narvekar, B. Peng, M. Leonetti, J. Sinapov, M. E. Taylor, and P. Stone, "Curriculum learning for reinforcement learning domains: A framework and survey," *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 7382–7431, 2020.