

COMPOSING TASK KNOWLEDGE WITH MODULAR SUCCESSOR FEATURE APPROXIMATORS

Wilka Carvalho^{*,1}

Angelos Filos²

Richard L. Lewis¹

Honglak lee^{1,3}

Satinder Singh¹

¹University of Michigan

²University of Oxford

³LG AI Research

ABSTRACT

Recently, the Successor Features and Generalized Policy Improvement (SF&GPI) framework has been proposed as a method for learning, composing, and transferring predictive knowledge and behavior. SF&GPI works by having an agent learn predictive representations (SFs) that can be combined for transfer to new tasks with GPI. However, to be effective this approach requires state features that are useful to predict, and these state-features are typically hand-designed. In this work, we present a novel neural network architecture, “Modular Successor Feature Approximators” (MSFA), where modules both discover what is useful to predict, and learn their own predictive representations. We show that MSFA is able to better generalize compared to baseline architectures for learning SFs and modular architectures for learning state representations.

1 INTRODUCTION

Consider a household robot that needs to learn tasks including picking up dirty dishes and cleaning up spills. Now consider that the robot is deployed and encounters a table with both a spill and a set of dirty dishes. Ideally this robot can combine its training behaviors to both clean up the spill and pickup the dirty dishes. We study this aspect of generalization: combining knowledge from multiple tasks.

Combining knowledge from multiple tasks is challenging because it is not clear how to synthesize either the behavioral policies or the value functions learned during training. This challenge is exacerbated when an agent also needs to generalize to novel appearances and environment configurations. Returning to our example, our robot might need to additionally generalize to both novel dirty dishes and to novel arrangements of chairs.

Successor features (SFs) and Generalized Policy Improvement (GPI) provide a mechanism to combine knowledge from multiple training tasks (Barreto et al., 2017; 2020). SFs are predictive representations that estimate how much state-features (known as “cumulants”) will be experienced given a behavior. By assuming that reward has a linear relationship between cumulants and a task vector, an agent can efficiently *compute* how much reward it can expect to obtain from a given behavior. If the agent knows multiple behaviors, it can leverage GPI to compute which behavior would provide the most reward (see Figure 2 for an example). However, SF&GPI commonly assume hand-designed cumulants and don’t have a mechanism for generalizing to novel environment configurations.

Modular architectures are a promising method for generalizing to distributions outside of the training distribution (Goyal et al., 2019; Madan et al., 2021). Recently, Carvalho et al. (2021a) presented “FARM” and showed that learning multiple state modules enabled generalization to environments with unseen environment parameters (e.g. to larger maps with more objects). In this work, we hypothesize that modules can further be leveraged to discover state-features that are useful to predict.

^{*}Contact author: wcarvalh@umich.edu.




Figure 1: (1) FARM learns multiple state modules. This promotes generalization to novel environments. However, it has no mechanism for combining task solutions. (2) USFA learns a single monolithic architecture for predicting SFs and can combine task solutions. However, it relies on hand-designed state features and has no mechanism for generalization to novel environments. (3) We combine the benefits of both. We leverage modules for reward-driven discovery of state features (SFs) that are useful to predict. These form the basis of their own predictive representations (SFs) and enables combining task solutions in novel environments.

We present “Modular Successor Feature Approximators” (MSFA), a novel neural network for discovering, composing, and transferring predictive knowledge and behavior via SF&GPI. MSFA is composed of a set of modules, which each learn their own state-features and corresponding predictive representations (SFs). **Our core contribution** is showing that an inductive bias for modularity can enable reward-driven discovery of state-features that are useful for zero-shot transfer with SF&GPI. We exemplify this with a simple state-feature discovery method presented in Barreto et al. (2018) where the dot-product between state-features and a task vector is regressed to environment reward. This method enabled transfer with SF&GPI in a continual learning setting but had limited success in the zero-shot transfer settings we study. While there are other methods for state-feature discovery, they add training complexity with mutual information objectives (Hansen et al., 2019) or meta-gradients (Veeriah et al., 2019). With MSFA, by adding *only an architectural bias for modularity*, we discover state-features that (1) support zero-shot transfer competitive with hand-designed features, and (2) enable zero-shot transfer in visually diverse, procedurally generated environments. We are hopeful that our architectural bias can be leveraged with other discovery methods in future work.

2 RELATED WORK ON GENERALIZATION IN RL

Hierarchical RL (HRL) is one dominant approach for combining task knowledge. The basic idea is that one can sequentially combine policies in time by having a “meta-policy” that sequentially activates “low-level” policies for protracted periods of time. By leveraging hand-designed or pre-trained low-level policies, one can generalize to longer instructions (Oh et al., 2017; Corona et al., 2020), to new instruction orders (Brooks et al., 2021), and to novel subtask graphs (Sohn et al., 2020; 2022). We differ in that we focus on combining policies *concurrently* in time as opposed to sequentially in time. To do so, we develop a modular neural network for the SF&GPI framework.

SFs are predictive representations that represent the current state as a summary of the *successive* features to follow (see §3 for a formal definition). By combining them with Generalized Policy Improvement, researchers have shown that they can transfer behaviors across object navigation tasks (Borsa et al., 2019; Zhang et al., 2017; Zhu et al., 2017), across continuous control tasks (Hunt et al., 2019), and within an HRL framework (Barreto et al., 2019). However, these works tend to require hand-designed cumulants which are cumbersome to design for every new environment. In our work, we integrate SFs with Modular RL to facilitate reward-driven discovery of cumulants and improve successor feature learning.

Modular RL (MRL) (Russell & Zimdars, 2003) is a framework for generalization by combining value functions. Early work dates back to (Singh, 1992), who had a mixture-of-experts system select between separately trained value functions. Since then, MRL has been applied to generalize across robotic morphologies (Huang et al., 2020), to novel task-robot combinations (Devin et al., 2017; Haarnoja et al., 2018), and to novel language

instructions (Logeswaran et al., 2021). MSFA, is the first to integrate MRL with SF&GPI. This integration enables combining task solutions in novel environment configurations.

Generalizing to novel environment configurations with modules. Goyal et al. (2019) showed that leveraging modules to learn a *state function* improved out-of-distribution generalization. Carvalho et al. (2021a) showed that a modified attention mechanism led to strong generalization improvements with RL. MSFA differs from both in that it employ modules for learning *value functions* in the form of SFs. This enables a principled way to compose task knowledge while additionally generalizing to novel environment configurations.

3 PROBLEM SETTING AND BACKGROUND

We study a reinforcement learning agent’s ability to transfer knowledge between tasks in an environment. During training, the experiences n_{train} tasks $\mathbb{M}_{\text{train}} = \{\mathcal{M}_i\}_{i=1}^{n_{\text{train}}}$, sampled from a training distribution $p_{\text{train}}(\mathcal{M})$. During testing, the agent is evaluated on n_{test} tasks, $\{\mathcal{M}_i\}_{i=1}^{n_{\text{test}}}$, sampled from a testing distribution $p_{\text{test}}(\mathcal{M})$. Each task \mathcal{M}_i is specified as Partially Observable Markov Decision Process (POMDP), $\mathcal{M}_i = \langle \mathcal{S}^e, \mathcal{A}, \mathcal{X}, R, p, f_x \rangle$. Here, \mathcal{S}^e , \mathcal{A} and \mathcal{X} are the environment state, action, and observation spaces. $p(\cdot | s_t^e, a_t)$ specifies the next-state distribution based on taking action a_t in state s_t^e , and $f_x(s_t^e)$ maps the underlying environment state to an observation x_t . We focus on tasks where rewards are parameterized by a task vector w , i.e. $r_t^w = R(s_t^e, a_t, s_{t+1}^e, w)$ is the reward obtained for transition (s_t^e, a_t, s_{t+1}^e) given task vector w . Since this is a POMDP, we need to learn a state function that maps histories to agent state representations. We do so with a recurrent function: $s_t = s_\theta(x_t, s_{t-1}, a_{t-1})$. Given this learned state, we want to obtain a behavioral policy $\pi(s_t)$ that best maximises the expected reward it will obtain when taking an action a_t at a state s_t : $Q_t^{\pi, w} = Q^{\pi, w}(s_t, a_t) = \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r_t^w]$.




Figure 2: **High-level diagram of how MSFA can be leveraged for transfer with SF&GPI.** During training, we can have the agent learn policies for tasks—e.g. “open drawer” and “open fridge”. Each task leads the agent to experience different aspects of the environment—e.g. a “fork” during “open drawer” or an “apple” during “open fridge”. We can leverage MSFA to have different modules learn different “cumulants”, ϕ , and SFs, ψ . For example, module 1 (θ_1) can estimate SFs for apple cumulants. Module SFs are combined to form the SF for a policy. When the agent wants to transfer its knowledge to a test task—e.g., “get milk”—it can compute Q-values for that task as a dot-product with the SFs of each training task. The highest Q-value is then used to select actions.

Transfer with SF&GPI. In order to leverage SFs (Barreto et al., 2017), one assumes an agent has access to state features known as “cumulants”, $\phi_t = \phi(s_t, a_t, s_{t+1})$. Given a

behavioral policy $\pi(a|s)$, SFs are a type of value function that use ϕ_t as pseudo-rewards:

$$\psi_t^\pi = \psi^\pi(s_t, a_t) = \mathbb{E}_\pi \left[\sum_{i=0}^{\infty} \gamma^i \phi_{t+i} \right] \quad (1)$$

If reward is (approximately) $r_t^w = \phi_t^\top w$, then action-values can be decomposed as $Q_t^{\pi,w} = \psi_t^{\pi^\top} w$. This is interesting because it provides an easy way to *reuse* task-agnostic features ψ_t^π for new tasks.

We can re-use the SFs we’ve learned from training tasks $\mathbb{M}_{\text{train}}$ for transfer with GPI. Assume we have learned (potentially optimal) policies $\{\pi_i\}_{i=1}^{n_{\text{train}}}$ and their corresponding SFs $\{\psi^{\pi_i}(s, a)\}_{i=1}^{n_{\text{train}}}$. Given a test task w_{test} , we can obtain a new policy with GPI in two steps: (1) compute Q-values using the training task SFs (2) select actions using the highest Q-value. This operation is summarized as follows:

$$\pi(s_t; w_{\text{test}}) \in \arg \max_{a \in \mathcal{A}} \max_{i \in \{1, \dots, n_{\text{train}}\}} \{Q_t^{\pi_i, w_{\text{test}}}\} = \arg \max_{a \in \mathcal{A}} \max_{i \in \{1, \dots, n_{\text{train}}\}} \{\psi_t^{\pi_i \top} w_{\text{test}}\} \quad (2)$$

This is useful because the GPI theorem states that π will perform as well as all of the training policies, i.e. that $Q^{\pi, w_{\text{test}}}(s, a) \geq \max_i Q^{\pi_i, w_i}(s, a) \forall (s, a) \in (\mathcal{S} \times \mathcal{A})$ (Barreto et al., 2017).

SF&GPI enable transfer by exploiting structure in the RL problem: a policy that maximizes a value function is guaranteed to perform at least as well as the policy that defined that value function. However, SF&GPI relies on combining a fixed set of SFs. Another form of transfer comes from “Universal Value Function Approximators” (UVFAs) (Schaul et al., 2015), which add the task-vector w as a parameter to a Q-approximator parameterized by θ , $Q_\theta(s, a, w)$. If Q_θ is smooth with respect to w , then Q_θ should generalize to test tasks nearby to train tasks in task space. Borsa et al. (2019) showed that one could combine the benefits of both with “Universal Successor Feature Approximators”. Since rewards r^w , and therefore task vectors w , reference deterministic task policies π_w , one can parameterize successor feature approximators with task-vectors $\tilde{\psi}^{\pi_w} = \tilde{\psi}^w \approx \psi_\theta(s, a, w)$. However, USFA assumed hand-designed cumulants. We introduce an architecture for reward-driven discovery of cumulants and improved function approximation of universal successor features.

4 MODULAR SUCCESSOR FEATURE APPROXIMATORS

We propose a new architecture *Modular Successor Feature Approximators* (MSFA) for approximating SFs, shown in Figure 3. Our hypothesis is that learning cumulants and SFs with modules improves zero-shot composition of task knowledge with SF&GPI. MSFA accomplishes this by learning n state modules $\{s_{\theta_k}\}_{k=1}^n$ that evolve with independent parameters θ_k and have sparse inter-module information flow. MSFA then produces modular cumulants $\{\tilde{\phi}_t^{(k)}\}_{k=1}^n$ and SFs $\{\tilde{\psi}_t^{\pi, k}\}_{k=1}^n$ by having their computations depend **only** on individual module-states. For example, a cumulant may correspond to information about apples, and would be a function **only** of the module representing state information related to apples. This is in contrast to prior work, which learns a single monolithic prediction module for computing cumulants and SFs (see Figure 3).

The rest of section is structured as follows. In section §4.1, we derive Modular Successor Feature Learning within the Modular RL framework. We then describe our architecture, MSFA, for learning modular successor features in §4.2. In §4.3, we describe how to generate behavior with MSFA. Finally, we describe the learning algorithm for MSFA in section §4.4.

4.1 MODULAR SUCCESSOR FEATURE LEARNING

Following the Modular RL framework (Russell & Zimdars, 2003), we assume that reward has an additive structure $R(s_t, a_t, s_{t+1}) = \sum_k R^k(s_t, a_t, s_{t+1}) = \sum_k R_t^{(k)}$, where $R_t^{(k)}$ is the reward of the k -th module. We enforce that every module decomposes reward into an inner-product between its own task-description $w^{(k)}$ and task-agnostic cumulants $\phi^{(k)} \in \mathbb{R}$: $R_t^{(k)} = \phi_t^{(k)} \cdot w^{(k)}$. Here, we simply break up the task vector into n pieces so individual

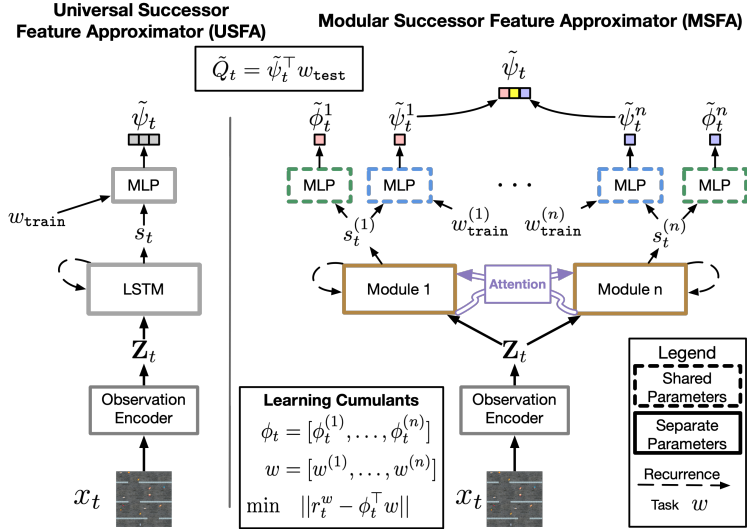


Figure 3: **Left: Universal Successor Feature Approximator (USFA)** learns a single, monolithic successor feature estimator that uses **hand-designed** cumulants. **Right: Modular Successor Feature Approximator (MSFA)** learns a set of successor feature modules, each with their own functions for (a) updating module-state, (b) computing cumulants, and (c) estimating successor features. Modules then share information with an attention mechanism. We hypothesize that isolated module computations facilitate learning cumulants that support generalization with GPI.

modules are responsible for subsets of the task vector. This allows us to decompose the action-value function as

$$Q^\pi(s_t, a_t, w) = \sum_{k=1}^n Q^{\pi,k}(s_t, a_t, w^{(k)}) = \sum_{k=1}^n \psi^{\pi,k}(s_t, a_t) \cdot w^{(k)} \quad (3)$$

where we now have *modular SFs* $\{\psi^{\pi,k}(s, a)\}_{k=1}^n$ (see the Appendix for a derivation). Rather than hand-designing modules or cumulants, we aim to discover them from the environment reward signal.

4.2 ARCHITECTURE

We learn a set of modules with states $\mathbb{S}_t = \{s_t^{(k)}\}_{k=1}^n$. They update at each time-step t with the observation x_t , the previous module-state $s_{t-1}^{(k)}$, and information from other modules $A_\theta(s_{t-1}^{(k)}, \mathbb{S}_{t-1})$. Following prior work (Santoro et al., 2018; Goyal et al., 2019; Carvalho et al., 2021a), we have A_θ combine transformer-style attention (Vaswani et al., 2017) with a gating mechanism (Parisotto et al., 2020) to enforce that inter-module interactions are sparse. Since A_θ is not the main contribution of this paper, we describe these computations in more detail with our notation in the Appendix. We summarize the high-level update below.

$$s_t^{(k)} = s_{\theta_k}(x_t, s_{t-1}^{(k)}, A_\theta(s_{t-1}^{(k)}, \mathbb{S}_{t-1})) \quad (4)$$

We learn modular cumulants and SFs by having sets of cumulants and SFs depend on individual module-states. Module cumulants depend on the module-state from the current and next time-step. Module SFs depend on the current module-state and on their subset of the task description. We summarize this below:

$$\tilde{\phi}_t^{(k)} = \phi_\theta(s_t^{(k)}, a_t, s_{t+1}^{(k)}) \quad \tilde{\psi}_t^{w,k} = \psi_\theta(s_t^{(k)}, a_t, w^{(k)}) \quad (5)$$

We highlight that cumulants share parameters but differ in their input. This suggests that the key is not having cumulants and SFs with separate parameters but that they are functions

of sparse subsets of state (rather than all state information). We show evidence for this hypothesis in Figure 6.

We concatenate module-specific cumulants and SFs to form the final outputs: $\tilde{\phi}_t = [\tilde{\phi}_t^{(1)}, \dots, \tilde{\phi}_t^{(n)}]$ and $\tilde{\psi}_t^w = \psi_\theta(s_t, a_t, w) = [\tilde{\psi}_t^{w,1}, \dots, \tilde{\psi}_t^{w,n}]$. Note that cumulants, are only used during learning, update with the module-state from the next time-step.

4.3 BEHAVIOR

During **training**, actions are selected in proportion to Q-values computed using task SFs as $\pi(s_t, w) \propto \tilde{Q}(s_t, a, w) = \psi_\theta(s_t, a, w)^\top w$. In practice we use an epsilon-greedy policy, though one can use other choices. During **testing**, we compute policies with GPI as $\pi(s_t, w_{\text{test}}) \in \arg \max_a \max_{z \in \mathbb{M}_{\text{train}}} \{\psi_\theta(s_t, a, z)^\top w_{\text{test}}\}$, where $\mathbb{M}_{\text{train}}$ are train task vectors.

4.4 LEARNING ALGORITHM

MSFA relies on three losses. The first loss, \mathcal{L}_Q , is a standard Q-learning loss, which MSFA uses to learn optimal policies for the training tasks. The main difference here is that MSFA uses a particular parameterization of the Q-function $Q^{\pi_w, w}(s, a) = \psi^{\pi_w}(s, a)^\top w$. The second loss, \mathcal{L}_ψ , is an SF learning loss, which we use as a regularizer to enforce that the Q-values follow the structure in the reward function $r_t^w = \phi_t^\top w$. For this, we again apply standard Q-learning but using SFs as value functions and cumulants as pseudo-rewards. The final loss, \mathcal{L}_ϕ is a loss for learning cumulants that grounds them in the environment reward signal. The losses are summarised as follows

$$\mathcal{L}_Q = \|r_t + \gamma \psi_\theta(s_{t+1}, a', w)^\top w - \psi_\theta(s_t, a_t, w)^\top w\|^2 \quad (6)$$

$$\mathcal{L}_\psi = \|\tilde{\phi}_t + \gamma \psi_\theta(s_{t+1}, a', w) - \psi_\theta(s_t, a_t, w)\|^2 \quad (7)$$

$$\mathcal{L}_\phi = \|r_t^w - \tilde{\phi}_t^\top w\|^2 \quad (8)$$

where $a' = \arg \max_a \psi_\theta(s_{t+1}, a, w)^\top w$. Selecting the next action via the combination of all modules ensures they individually convergence to optimal values (Russell & Zimdars, 2003). The final loss is $\mathcal{L} = \alpha_Q \mathcal{L}_Q + \alpha_\psi \mathcal{L}_\psi + \alpha_\phi \mathcal{L}_\phi$.

5 EXPERIMENTS




Figure 4: **We study an agent’s ability to combine task knowledge in three environments.** (a) In BabyAI, an agent learns to pick up one object type at a time during training. During testing, the agent must pickup combinations of object types while avoiding other object types. This is the setting used by USFA which assumed **hand-designed** cumulants. (b) In Procgen, we study extending this form of generalization to a visually diverse, procedurally generated environment. (c) In Minihack, we go beyond combining object navigation skills. Here, an agents needs to combine (1) avoiding teleportation traps, (2) avoiding monsters, and (3) partial visibility around the agent.

We study generalization when training behaviors must be combined concurrently in time in the presence of novel object appearances and layouts. The need to combine training behaviors

tests how well MSFA can leverage SF&GPI. Generalization to novel object appearances and layouts tests how well MSFA’s modular construction supports generalization to novel environment configurations.

Baselines. (1) **Universal Value Function Approximator (UVFA)** (Schaul et al., 2015), which takes the task as input: $Q_\theta(s, w)$. This comparison shows the transfer benefits of SF&GPI. (2) **UVFA with Feature-Attending Recurrent Modules (UVFA+FARM)** instead takes state-factors as input $Q_\theta(s^{(1)}, \dots, s^{(n)}, w)$. Each state-factor $s^{(k)}$ is the output of a FARM module. (3) **Modular Value Function Approximator(MVFA)** is an adaptation of (Haarnoja et al., 2018) where modules learn individual Q-values $Q_\theta^{(i)}(s^{(i)}, w^{(i)})$. Comparing to UVFA+FARM and MVFA enables us to study the benefits of leveraging modules for learning value functions in the form of SFs. (4) **Universal Successor Function Approximator (USFA)** (Borsa et al., 2019) leverages a single monolithic function for successor features with **hand-designed cumulants**. USFA is an upper-bound baseline that allows us to test the quality of cumulants and successor features that MSFA learns. We also test a variant of USFA with learned cumulants, **USFA-Learned- ϕ** , which shows how the architecture degrades without oracle cumulants.

Implementation. We implement the state modules of MSFA with FARM modules (Carvalho et al., 2021b). For UVFA and USFA, we learn a state representation with an LSTM (Hochreiter & Schmidhuber, 1997). We implement all ϕ , ψ , and Q functions with Mutli-layer Perceptrons. We train UVFA and UVFA+FARM with n-step Q-learning (Watkins & Dayan, 1992). When learning cumulants, USFA and MSFA have the exact same losses and learning algorithm. They both learn Q-values and SFs with n-step Q-learning. We use $n = 5$. When not learning cumulants, following (Borsa et al., 2019), USFA only learns SFs with n-step Q-learning. All agents are built with JAX (Bradbury et al., 2018) using the open-source ACME codebase (Hoffman et al., 2020) for reinforcement learning.

5.1 COMBINING OBJECT NAVIGATION TASK KNOWLEDGE




Figure 5: **MSFA matches USFA, which has hand-designed cumulants.** We show mean and standard error generalization episode return across 10 runs. We put a task’s L2 distance to the closest train task in parenthesis. USFA best generalizes to novel combinations of picking up and avoiding objects. Once USFA learns cumulants, its performance degrades significantly. UVFA-based methods struggle as more objects should be avoided or tasks are further in distance to train tasks.

MSFA learns modular functions for computing ϕ and estimating ψ . We hypothesize that this facilitates learning cumulants that respond to different aspects of the environment (e.g. to different object categories). This leads to the following research questions. **R1:** Can we recover prior generalization results that relied on hand-designed cumulants for different object categories? **R2:** How important is it to learn modular functions for ϕ and ψ ? **R3:** Without GPI, do learning modular functions for ϕ and ψ still aid in generalization?

Setup. We implement a simplified version of the object navigation task of (Borsa et al., 2019) in the (Chevalier-Boisvert et al., 2019) BabyAI environment. The environment consists of 3 instances of 4 object categories. **Observations** are partial and egocentric. **Actions:** the agent can rotate left or right, move forward, or pickup an object. When it picks up an object, following (Borsa et al., 2019), the object is respawned somewhere on the grid. **Task** vectors lie in $w \in \mathbb{R}^4$ with training tasks being the standard unit vectors. For example, $[0, 1, 0, 0]$ specifies the agent must obtain objects of category 2. **Generalization** tasks are linear combinations of training tasks. For example, $[-1, 1, -1, 1]$ specifies the agent must obtain categories 2 and 4 while *avoiding* categories 1 and 3. Borsa et al. (2019) showed that USFA could generalize with *hand-designed* cumulants that described whether an object was picked up. We describe challenges for this task in detail in the Appendix.

R1: MSFA is competitive with USFA, which uses oracle ϕ . Figure 5 shows USFA with a similar generalization trend to (Borsa et al., 2019). Tasks get more challenging as they are further from train tasks or involve avoiding more objects. For simply going to combinations of objects, USFA-Learned- ϕ does slightly worse than MSFA. However, with more objects to avoid, all methods except MSFA (including USFA-Learned- ϕ) degrades significantly. For comparison, we show performance by an oracle bread-first-search policy with access to ground-truth state (**BFS State Oracle**). All methods have room for improvement when objects must be avoided. In the appendix, we present heat-maps for how often object categories were picked up during different tasks. We find that MSFA most matches USFA, while USFA-Learned- ϕ commonly picks up all objects regardless of task.




Figure 6: **Learning modular ϕ and ψ is key to generalization and improves generalization even without GPI.** We show mean and standard error generalization episode return across 10 runs. (a) We ablate having modular functions for ϕ_θ and ψ_θ . Generalization results degrade significantly. (b) We ablate leveraging GPI for generalization from all SF-based methods. MSFA without GPI can outperform both USFA-Learned- ϕ with GPI and USFA without GPI. This shows the utility of modularity for generalization.

R2: Learning modular ϕ and ψ functions is critical for generalization. Learning an entangled function corresponds to learning a monolithic function for ψ or ϕ where we concatenate module-states, e.g. $\tilde{\psi}_t^w = \psi_\theta(s_t^{(1)}, \dots, s_t^{(n)}, a, w)$. Modular functions correspond to equation 5. Figure 6 (a) shows that without modular functions for ϕ and ψ , performance severely degrades. This also highlights that a naive combination of USFA+FARM—with entangled functions for ϕ and ψ —does not recover our generalization performance.

R3: Modularity alone improves generalization. For all SF-based methods, we remove GPI and select actions with a greedy policy: $\pi(s) = \arg \max_a \psi_\theta(s_t, a, w)^\top w$. Figure 6 (b) shows that GPI is critical for generalization with USFA as expected. USFA-Learned- ϕ benefits less from GPI (presumably because of challenges in learning ϕ). Interestingly, MSFA can generalize relatively well without GPI, sometimes doing better than USFA without GPI.

5.2 COMBINING OBJECT NAVIGATION TASK KNOWLEDGE WITH NOVEL APPEARANCES AND ENVIRONMENT CONFIGURATIONS

Beyond generalizing to combinations of tasks, agents will need to generalize to different layouts and appearances of objects. **R4:** Can MSFA enable combining task knowledge in a visually diverse, procedurally generated environment?

Setup. We leverage the “Fruitbot” environment within ProcGen (Cobbe et al., 2020). Here, an agent controls a paddle that tries to obtain certain categories of objects while avoiding

others. When the agent hits a wall or fence, it dies and the episode terminates. If the agent collects a non-task object, nothing happens. **Observations** are partial. **Actions**: At each time-step the agent moves one step forward and can move left or right or shoot pellets to open fences. **Training and generalization tasks** follow the same setup as §5.1.




Figure 7: **MSFA is able to combine task knowledge in a visually diverse, procedurally generated ProcGen environment.** We find that no method is able to do well when there are objects to avoid ($w = [1, 0, 0, -1]$) in this setting (see text for more). However, as more objects need to be collected MSFA best generalizes (10 runs).

R4: MSFA enables combining task knowledge of object navigation tasks in a visually diverse, procedurally generated environment. Figure 7 shows that when an agent has to generalize to collecting more objects, modular architectures generalize best, with MSFA doing best. When objects have to be avoided, we see that no architecture does well, though MSFA tends to do better. We observe that avoiding objects leads agents to hit walls. Since the agent always moves forward at each time-step in Fruibot, this makes avoiding objects a particularly challenging type of generalization.

6 DISCUSSION AND CONCLUSION

We have presented “Modular Successor Feature Approximators”, a modular neural network for learning cumulants and SFs produced by their own modules. We first showed that MSFA is competitive with prior object navigation generalization results that relied on hand-designed cumulants (§5.1). Afterwards, we showed that MSFA improves an agent’s ability to combine task knowledge in a visually diverse, procedurally generated environment (§5.2). We also show that MSFA can combine solutions of heterogeneous tasks (§B). Our ablations show that learning modular cumulants and SFs is critical for generalization with GPI.

We compared MSFA to (1) USFA, a monolithic architecture for learning cumulants and SFs; (2) FARM, an architecture which learns multiple state-modules but combines them with a monolithic Q-value function. Our results show that when learning cumulants, MSFA improves generalization with SF&GPI compared to USFA. Additionally, without GPI, MSFA as an architecture improves generalization as compared to both FARM and USFA.

Limitations. While we demonstrated reward-driven discovery of cumulants for transfer with SF&GPI, we focused on relatively simple task encodings. Future work can extend this to more expressive encodings such as language embeddings. Another limitation is that we did not explore more sophisticated state-feature discovery methods such as meta-gradients (Veeriah et al., 2019). Nonetheless, we think that MSFA provides an important insight for future work: that modularity is a simple but powerful inductive bias for discovering state-features that enable zero-shot transfer with SF&GPI.

Future directions. SFs are useful for exploration (Janz et al., 2019; Machado et al., 2020); for discovering and combining options (Barreto et al., 2019; Hansen et al., 2019); for transferring policies across environments (Zhang et al., 2017); for improving importance sampling (Fujimoto et al., 2021); and for learning policies from other agents (Filos et al., 2021). We hope that future work can leverage MSFA for improved state-feature discovery and SF-learning in all of these settings.

ACKNOWLEDGMENTS

This work was supported in part by a University of Michigan Rackham Merit Fellowship and grant from LG AI Research. The authors would also like to thank Andrew Lampinen, Cameron Allen, and the anonymous reviewers for their helpful comments.

REFERENCES

- André Barreto, Will Dabney, Rémi Munos, Jonathan J Hunt, Tom Schaul, Hado P van Hasselt, and David Silver. Successor features for transfer in reinforcement learning. *NIPS*, 30, 2017.
- André Barreto, Diana Borsa, John Quan, Tom Schaul, David Silver, Matteo Hessel, Daniel Mankowitz, Augustin Zidek, and Remi Munos. Transfer in deep reinforcement learning using successor features and generalised policy improvement. In *ICML*, pp. 501–510. PMLR, 2018.
- André Barreto, Diana Borsa, Shaobo Hou, Gheorghe Comanici, Eser Aygün, Philippe Hamel, Daniel Toyama, Shibl Mourad, David Silver, Doina Precup, et al. The option keyboard: Combining skills in reinforcement learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- André Barreto, Shaobo Hou, Diana Borsa, David Silver, and Doina Precup. Fast reinforcement learning with generalized policy updates. *PNAS*, 117(48):30079–30087, 2020.
- Diana Borsa, André Barreto, John Quan, Daniel Mankowitz, Rémi Munos, Hado Van Hasselt, David Silver, and Tom Schaul. Universal successor features approximators. *ICLR*, 2019.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Ethan Brooks, Janarthanan Rajendran, Richard L Lewis, and Satinder Singh. Reinforcement learning of implicit and explicit control flow instructions. In *International Conference on Machine Learning*, pp. 1082–1091. PMLR, 2021.
- Wilka Carvalho, Andrew Lampinen, Kyriacos Nikiforou, Felix Hill, and Murray Shanahan. Feature-attending recurrent modules for generalization in reinforcement learning. *arXiv*, 2021a.
- Wilka Carvalho, Anthony Liang, Kimin Lee, Sungryull Sohn, Honglak Lee, Richard L Lewis, and Satinder Singh. Reinforcement learning for sparse-reward object-interaction tasks in a first-person simulated 3d environment. *IJCAI*, 2021b.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. Babyai: A platform to study the sample efficiency of grounded language learning. *arXiv preprint arXiv:1810.08272*, 2018.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. BabyAI: First steps towards grounded language learning with a human in the loop. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJeXCo0cYX>.
- Karl Cobbe, Chris Hesse, Jacob Hilton, and John Schulman. Leveraging procedural generation to benchmark reinforcement learning. In *International conference on machine learning*, pp. 2048–2056. PMLR, 2020.
- Rodolfo Corona, Daniel Fried, Coline Devin, Dan Klein, and Trevor Darrell. Modular networks for compositional instruction following. *arXiv preprint arXiv:2010.12764*, 2020.

- Coline Devin, Abhishek Gupta, Trevor Darrell, Pieter Abbeel, and Sergey Levine. Learning modular neural network policies for multi-task and multi-robot transfer. In *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 2169–2176. IEEE, 2017.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Vlad Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures. In *International conference on machine learning*, pp. 1407–1416. PMLR, 2018.
- Angelos Filos, Clare Lyle, Yarín Gal, Sergey Levine, Natasha Jaques, and Gregory Farquhar. Psiphi-learning: Reinforcement learning with demonstrations using successor features and inverse temporal difference learning. In *International Conference on Machine Learning*, pp. 3305–3317. PMLR, 2021.
- Scott Fujimoto, David Meger, and Doina Precup. A deep reinforcement learning approach to marginalized importance sampling with the successor representation. In *International Conference on Machine Learning*, pp. 3518–3529. PMLR, 2021.
- Anirudh Goyal, Alex Lamb, Jordan Hoffmann, Shagun Sodhani, Sergey Levine, Yoshua Bengio, and Bernhard Schölkopf. Recurrent independent mechanisms. *arXiv preprint arXiv:1909.10893*, 2019.
- Tuomas Haarnoja, Vitchyr Pong, Aurick Zhou, Murtaza Dalal, Pieter Abbeel, and Sergey Levine. Composable deep reinforcement learning for robotic manipulation. In *2018 IEEE international conference on robotics and automation (ICRA)*, pp. 6244–6251. IEEE, 2018.
- Steven Hansen, Will Dabney, Andre Barreto, Tom Van de Wiele, David Warde-Farley, and Volodymyr Mnih. Fast task inference with variational intrinsic successor features. *arXiv preprint arXiv:1906.05030*, 2019.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Matt Hoffman, Bobak Shahriari, John Aslanides, Gabriel Barth-Maron, Feryal Behbahani, Tamara Norman, Abbas Abdolmaleki, Albin Cassirer, Fan Yang, Kate Baumli, et al. Acme: A research framework for distributed reinforcement learning. *arXiv preprint arXiv:2006.00979*, 2020.
- Wenlong Huang, Igor Mordatch, and Deepak Pathak. One policy to control them all: Shared modular policies for agent-agnostic control. In *International Conference on Machine Learning*, pp. 4455–4464. PMLR, 2020.
- Jonathan Hunt, Andre Barreto, Timothy Lillicrap, and Nicolas Heess. Composing entropic policies using divergence correction. In *International Conference on Machine Learning*, pp. 2911–2920. PMLR, 2019.
- David Janz, Jiri Hron, Przemysław Mazur, Katja Hofmann, José Miguel Hernández-Lobato, and Sebastian Tschiatschek. Successor uncertainties: exploration and uncertainty in temporal difference learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Steven Kapturowski, Georg Ostrovski, John Quan, Remi Munos, and Will Dabney. Recurrent experience replay in distributed reinforcement learning. In *International conference on learning representations*, 2018.
- Lajanugen Logeswaran, Wilka Torrico Carvalho, and Honglak Lee. Learning compositional tasks from language instructions. In *Deep RL Workshop NeurIPS 2021*, 2021.
- Marlos C Machado, Marc G Bellemare, and Michael Bowling. Count-based exploration with the successor representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, number 04, pp. 5125–5133, 2020.

- Kanika Madan, Nan Rosemary Ke, Anirudh Goyal, Bernhard Schölkopf, and Yoshua Bengio. Fast and slow learning of recurrent independent mechanisms. *arXiv preprint arXiv:2105.08710*, 2021.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Junhyuk Oh, Satinder Singh, Honglak Lee, and Pushmeet Kohli. Zero-shot task generalization with multi-task deep reinforcement learning. In *International Conference on Machine Learning*, pp. 2661–2670. PMLR, 2017.
- Emilio Parisotto, Francis Song, Jack Rae, Razvan Pascanu, Caglar Gulcehre, Siddhant Jayakumar, Max Jaderberg, Raphael Lopez Kaufman, Aidan Clark, Seb Noury, et al. Stabilizing transformers for reinforcement learning. In *International conference on machine learning*, pp. 7487–7498. PMLR, 2020.
- Stuart J Russell and Andrew Zimdars. Q-decomposition for reinforcement learning agents. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 656–663, 2003.
- Mikayel Samvelyan, Robert Kirk, Vitaly Kurin, Jack Parker-Holder, Minqi Jiang, Eric Hambro, Fabio Petroni, Heinrich Küttler, Edward Grefenstette, and Tim Rocktäschel. Minihack the planet: A sandbox for open-ended reinforcement learning research. *arXiv preprint arXiv:2109.13202*, 2021.
- Adam Santoro, Ryan Faulkner, David Raposo, Jack Rae, Mike Chrzanowski, Theophane Weber, Daan Wierstra, Oriol Vinyals, Razvan Pascanu, and Timothy Lillicrap. Relational recurrent neural networks. *Advances in neural information processing systems*, 31, 2018.
- Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In *International conference on machine learning*, pp. 1312–1320. PMLR, 2015.
- Satinder Pal Singh. Transfer of learning by composing solutions of elemental sequential tasks. *Machine Learning*, 8(3):323–339, 1992.
- Sungryull Sohn, Hyunjae Woo, Jongwook Choi, and Honglak Lee. Meta reinforcement learning with autonomous inference of subtask dependencies. *arXiv preprint arXiv:2001.00248*, 2020.
- Sungryull Sohn, Hyunjae Woo, Jongwook Choi, Lyubing Qiang, Izzeddin Gur, Aleksandra Faust, and Honglak Lee. Fast inference and transfer of compositional task structures for few-shot task generalization. In *Uncertainty in Artificial Intelligence*, pp. 1857–1865. PMLR, 2022.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv*, 2017.
- Vivek Veeriah, Matteo Hessel, Zhongwen Xu, Janarthanan Rajendran, Richard L Lewis, Junhyuk Oh, Hado P van Hasselt, David Silver, and Satinder Singh. Discovery of useful questions as auxiliary tasks. *Advances in Neural Information Processing Systems*, 32, 2019.
- Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3):279–292, 1992.
- Jingwei Zhang, Jost Tobias Springenberg, Joschka Boedecker, and Wolfram Burgard. Deep reinforcement learning with successor features for navigation across similar environments. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2371–2378. IEEE, 2017.
- Yuke Zhu, Daniel Gordon, Eric Kolve, Dieter Fox, Li Fei-Fei, Abhinav Gupta, Roozbeh Mottaghi, and Ali Farhadi. Visual semantic planning using deep successor representations. In *Proceedings of the IEEE international conference on computer vision*, pp. 483–492, 2017.