

Context-Specific Representation Abstraction for Deep Option Learning

Marwa Abdulhai^{1,2}, Dong-Ki Kim^{1,2}, Matthew Riemer^{2,3}, Miao Liu^{2,3},
Gerald Tesauro^{2,3}, Jonathan P. How^{1,2}

¹ MIT LIDS

² MIT-IBM Watson AI Lab

³ IBM Research

abdulhai@mit.edu, dkkim93@mit.edu, mdriemer@us.ibm.com, miao.liu1@ibm.com,
gtesauro@us.ibm.com, jhow@mit.edu

Abstract

Hierarchical reinforcement learning has focused on discovering temporally extended actions, such as options, that can provide benefits in problems requiring extensive exploration. One promising approach that learns these options end-to-end is the option-critic (OC) framework. We examine and show in this paper that OC does not decompose a problem into simpler sub-problems, but instead increases the size of the search over policy space with each option considering the entire state space during learning. This issue can result in practical limitations of this method, including sample inefficient learning. To address this problem, we introduce Context-Specific Representation Abstraction for Deep Option Learning (CRADOL), a new framework that considers both temporal abstraction and context-specific representation abstraction to effectively reduce the size of the search over policy space. Specifically, our method learns a factored belief state representation that enables each option to learn a policy over only a subsection of the state space. We test our method against hierarchical, non-hierarchical, and modular recurrent neural network baselines, demonstrating significant sample efficiency improvements in challenging partially observable environments.

Introduction

Hierarchical reinforcement learning (HRL) provides a principled framework for decomposing problems into natural hierarchical structures (Flet-Berliac 2019). By factoring a complex task into simpler sub-tasks, HRL can offer benefits over non-HRL approaches in solving challenging large-scale problems efficiently. Much of HRL research has focused on discovering temporally extended high-level actions, such as options, that can be used to improve learning and planning efficiency (Sutton, Precup, and Singh 1999; Precup and Sutton 2000). Notably, the option-critic (OC) framework has become popular because it can learn deep hierarchies of both option policies and termination conditions from scratch without domain-specific knowledge (Bacon, Harb, and Precup 2017; Riemer, Liu, and Tesauro 2018). However, while OC provides a theoretically grounded framework for HRL, it is known to exhibit common failure cases in practice including lack of option diversity, short option duration, and large sample complexity (Kamat and Precup 2020).

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Our key insight in this work is that OC suffers from these problems at least partly because it considers the entire state space during option learning and thus fails to reduce problem complexity as intended. For example, consider the partially observable maze domain in Figure 1. The objective in this domain is to pick up the key and unlock the door to reach the green goal. A non-HRL learning problem in this setting can be viewed as a search over the space of deterministic policies: $|\mathcal{A}|^{|\mathcal{B}|}$, where $|\mathcal{A}|$ and $|\mathcal{B}|$ denote the size of the action space and belief state space, respectively. Frameworks such as OC consider the entire state space or belief state when learning options, which can be viewed themselves as policies over the entire state space, naively resulting in an increased learning problem size of $|\Omega|(|\mathcal{A}|^{|\mathcal{B}|})$, where $|\Omega|$ denotes the number of options. Note that this has not yet considered the extra needed complexity for learning policies to select over and switch options. Unfortunately, because $|\Omega| \geq 1$, learning options in this way can then only increase the effective complexity of the learning problem.

To address the problem complexity issue, we propose to consider both temporal abstraction, in the form of high-level skills, and context-specific abstraction over factored belief states, in the form of latent representations, to learn options that can fully benefit from hierarchical learning. Considering the maze example again, the agent does not need to consider the states going to the goal when trying to get the key (see Option 1 in Figure 1). Similarly, the agent only needs to consider states that are relevant to solving a specific sub-task. Hence, this context-specific abstraction leads to the desired reduction in problem complexity, and an agent can mutually benefit from having both context-specific and temporal abstraction for hierarchical learning.

Remark 1 *With context-specific abstraction for each option, the learning problem size can be reduced such that:*

$$\sum_{\omega \in \Omega} |\mathcal{A}|^{|\mathcal{B}_\omega|} < |\mathcal{A}|^{|\mathcal{B}|} < |\Omega|(|\mathcal{A}|^{|\mathcal{B}|}), \quad (1)$$

where $\mathcal{B}_\omega \subset \mathcal{B}$ denotes the subset of the belief state space for each option (or sub-policy) ω and $|\mathcal{B}_\omega| < |\mathcal{B}|$. We further illustrate this idea in the tree diagram of Figure 1, where each option maps to only a subset of the state space.

Contribution. With this insight, we introduce a new option learning framework: Context-Specific Representation




Figure 1: Temporal abstraction in this maze consists of 3 options; getting the key, opening the door, and going to the green goal. With context-specific belief state abstraction, the agent should only focus on different factors of the belief state pertaining to the highlighted sub-task. Equipped with both temporal and context-specific abstraction, our method achieves a significant reduction in the learning problem size compared to considering the entire belief state space for each option as highlighted in Remark 1.

Abstraction for Deep Option Learning (CRADOL).

CRADOL provides options with the flexibility to select over dynamically learned abstracted components. This allows effective decomposition of problem size when scaling up to large state and action spaces, leading to faster learning. Our experiments demonstrate how CRADOL can improve performance over key baselines in partially observable settings.

Problem Setting and Notation

Partially Observable Markov Decision Process. An agent’s interaction in the environment can be represented by a partially observable Markov decision process (POMDP), formally defined as a tuple $\langle \mathcal{S}, \mathcal{A}, P, \mathcal{R}, \mathcal{X}, \mathcal{O}, \gamma \rangle$ (Kaelbling, Littman, and Cassandra 1998). \mathcal{S} is the state space, \mathcal{A} is the action space, P is the state transition probability function, \mathcal{R} is the reward function, \mathcal{X} is the observation space, \mathcal{O} is the observation probability function, and $\gamma \in [0, 1)$ is the discount factor. An agent executes an action a according to its stochastic policy $a \sim \pi_\theta(a|b)$ parameterized by $\theta \in \Theta$, where $b \in \mathcal{B}$ denotes the belief state and the agent uses the observation history up to the current timestep $x_{0:t}$ to form its belief b_t . Each observation $x \in \mathcal{X}$ is generated according to $x \sim \mathcal{O}(s)$. An action a yields a transition from the current state $s \in \mathcal{S}$ to the next state $s' \in \mathcal{S}$ with probability $P(s'|s, a)$, and an agent obtains a reward $r \sim \mathcal{R}(s, a)$. An agent’s goal is to find the deterministic optimal policy π^* that maximizes its expected discounted return $\mathbb{E}_{\pi_\theta} [\sum_{t=0}^{\infty} \gamma^t r_t] \forall \theta \in \Theta$.

Option-Critic Framework. A Markovian option $\omega \in \Omega$ consists of a triple $\langle I_\omega, \pi_\omega, \beta_\omega \rangle$. $I_\omega \in \mathcal{S}$ is the initiation set, π_ω is the intra-option policy, and $\beta_\omega: \mathcal{S} \rightarrow [0, 1]$ is the option termination function (Sutton, Precup, and Singh 1999). Similar to option discovery methods (Mankowitz, Mann, and Mannor 2016; Daniel et al. 2016; Vezhnevets et al. 2016), we also assume that all options are available in each state. MDPs with options become semi-MDPs (Puterman 1994) with an associated optimal value function over options $V_\Omega^*(s)$ and option-value function $Q_\Omega^*(s, \omega)$. Bacon, Harb, and Precup (2017) introduces a gradient-based framework for learning sub-policies represented by an option. An option ω is selected according to a policy over options $\pi_\Omega(\omega|s)$, and $\pi_\omega(a|s)$ selects a primitive action a until the termination of this option (as determined by $\beta_\omega(s)$), which triggers a repetition of this procedure. The framework also models another

auxiliary value function $Q_U(s, \omega, a)$ for computing the gradient of π_ω .

Factored POMDP. In this work, we further assume that the environment structure can be modeled as a factored MDP in which the state space \mathcal{S} is factored into n variables: $\mathcal{S} = \{\mathcal{S}^1, \dots, \mathcal{S}^n\}$, where each \mathcal{S}^i takes on values in a finite domain and $i \in \{1, \dots, n\}$ (Guestrin et al. 2003). Under this assumption, a state $s \in \mathcal{S}$ is an assignment of values to the set of state variables: $s = \{s^1, \dots, s^n\}$ such that $\mathcal{S} \subseteq \mathcal{S}^1 \times \dots \times \mathcal{S}^n$. These variables are also assumed to be largely independent of each other with a small number of causal parents $\mathcal{P}^{\mathcal{S}^i}$ to consider at each timestep, resulting in the following simplification of the transition function:

$$P(s'|s, a) \approx \prod_{i=1}^n P^i(s'^i | \mathcal{P}^{\mathcal{S}^i}(s), a). \quad (2)$$

This implies that in POMDP settings we can match the factored structure of the state space by representing our belief state in a factored form as well with $\mathcal{B} = \{\mathcal{B}^1, \dots, \mathcal{B}^n\}$ and $b = \{b^1, \dots, b^n\}$. Another consequence of the factored MDP assumption is context-specific independence.

Context-Specific Independence. At any given timestep, only a small subset of the factored state space may be necessary for an agent to be aware of. As per Bouillier et al. (1996), we define this subset to be a *context*. Formally, a context Z is a pair (Z, \mathcal{Z}) , where $Z \subseteq \mathcal{S}$ is a subset of state space variables and \mathcal{Z} is the space of possible joint assignments of state space variables in the subset. A state s is in the context (Z, \mathcal{Z}) when its joint assignment of variables in Z is present in \mathcal{Z} . Two variables $Y, Y' \subseteq \mathcal{S} \setminus Z$ are contextually independent under (Z, \mathcal{Z}) if $Pr(Y|Y', Z = z) = Pr(Y|Z = z) \forall z \in \mathcal{Z}$. This independence relation is referred to as context-specific independence (CSI) (Chitnis et al. 2020). CSI in the state space also implies that this same structure can be leveraged in the belief state for POMDP problems.

Context-Specific Belief State Abstractions for Option-Critic Learning

We now outline how factored state space structure and context specific independence can be utilized by OC to decrease the size of the search problem over policy space. We first note that the size of the search problem over the policy space

of OC $|\pi_{\text{OC}}|$ can be decomposed into the sum of the search problem sizes for each sub-policy within the architecture:

$$|\pi_{\text{OC}}| = |\pi_{\Omega}| + \sum_{\omega \in \Omega} |\pi_{\omega}| + |\beta_{\omega}|. \quad (3)$$

OC can provide a decrease in the size of the learning problem over the flat RL method when $|\pi_{\text{OC}}| < |\mathcal{A}|^{|\mathcal{B}|}$. This decrease is possible if an agent leverages context-specific independence relationships with respect to each option ω such that only a subset of the belief state variables are considered $\mathcal{B}_{\omega} \subset \mathcal{B}$, where $\mathcal{B}_{\omega} = \{\mathcal{B}^1, \dots, \mathcal{B}^k\}$ and $k < n$, implying that $|\mathcal{B}_{\omega}| \ll |\mathcal{B}|$ because the belief state space gets exponentially smaller with each removed variable. The abstracted belief state $b_{\omega} \in \mathcal{B}_{\omega}$ is then sent as an input to the intra-option policy $\pi_{\omega}(a|b_{\omega})$ and termination policy $\beta_{\omega}(b_{\omega})$, dictating the size of the learning problem for each:

$$|\pi_{\omega}| = |\mathcal{A}|^{|\mathcal{B}_{\omega}|}, \quad |\beta_{\omega}| = 2^{|\mathcal{B}_{\omega}|}. \quad (4)$$

However, a challenge becomes that we must also consider the size of the learning problem for the policy over options $|\pi_{\Omega}|$. π_{Ω} cannot simply focus on an abstracted belief state in the context of an individual option ω because it must reason about the impact of all options. Naively sending the full belief state b to π_{Ω} is unlikely to make the learning problem smaller for a large class of problems because we cannot achieve gains over the flat policy in this case if $|\Omega|$ is of comparable or greater size than $|\mathcal{A}|$. In this work, we address this issue by sending the policy over options its own compact abstracted belief state $b_{\Omega} \in \mathcal{B}_{\Omega}$, where $|\mathcal{B}_{\Omega}| < |\mathcal{B}|$. As a result, the learning problem size of $\pi_{\Omega}(\omega|b_{\Omega})$ is $|\Omega|^{|\mathcal{B}_{\Omega}|}$ and the total problem size for $|\pi_{\text{OC}}|$ is:

$$|\pi_{\text{OC}}| = |\Omega|^{|\mathcal{B}_{\Omega}|} + \sum_{\omega \in \Omega} |\mathcal{A}|^{|\mathcal{B}_{\omega}|} + 2^{|\mathcal{B}_{\omega}|}. \quad (5)$$

Note that b_{Ω} represents a different kind of belief state abstraction than b_{ω} in that it must consider all factors of the state space. b_{Ω} , however, should also consider much less detail than is contained in b for each factor, because π_{Ω} is only concerned with the higher level semi-MDP problem that operates at a much slower time scale when options last for significant lengths before terminating and does not need to reason about the intricate details needed to decide on individual actions at each timestep. As a result, we consider settings where $|\mathcal{B}_{\Omega}| \ll |\mathcal{B}|$ and $|\mathcal{B}_{\omega}| \ll |\mathcal{B}|$, ensuring that an agent leveraging our CRADOL framework solves a smaller learning problem than either flat RL or OC as traditionally applied. Finally, the auxiliary value function Q_U , which is used to compute the gradient of π_{ω} , must consider rewards accumulated across options similarly to π_{Ω} . Thus, we also send its own compact abstracted belief state $b_U \in \mathcal{B}_U$ as an input to Q_U such that the resulting problem size is small $|\mathcal{B}_U| \ll |\mathcal{B}|$.

Learning Context-Specific Abstractions with Recurrent Independent Mechanisms

In this section, we present a new framework, CRADOL, based on our theoretical motivation outlined for applying abstractions over belief state representations in the last section.

In particular, our approach leverages recurrent independent mechanisms (Goyal et al. 2021) to model factored context-specific belief state abstractions for deep option learning.

Algorithm Overview

Figure 2 shows a high-level overview of CRADOL. CRADOL uses small LSTM networks to provide compact abstract representations for b_{Ω} and b_U . It additionally leverages a set of mechanisms to model the factored belief state b while only sending a subset of this representation b_{ω} to each option during its operation.

Option Learning

Our framework for option learning broadly consists of four modules: option attention, a recurrent layer, sparse communication, and fully connected (see Figure 2). We choose to model each option based on a group of K mechanisms because these mechanisms can learn to specialize into sub-components following the general structure assumed in the environment for factored POMDPs. This in turn allows our learned options to self-organize into groupings of belief state components.

Option Attention Module. The option attention module chooses which mechanisms out of K mechanisms (M_1, \dots, M_K) are activated for $\omega \in \Omega$ by passing the option through a look-up table $W^{\text{table}} \in \mathbb{R}^{|\Omega| \times K}$. The lookup table ensures a fixed context selection over time (i.e. a fixed mapping to a subset of the belief state space) by an option as considering many different subsets with the same option would be effectively equivalent to considering the entire belief state space. This would lead to a large search for policy space as previously done by OC. To yield the output of the attention module h_{ω}^{attn} , we apply the following:

$$h_{\omega}^{\text{attn}} = \text{softmax}(W_{\omega}^{\text{table}})^T x W^{\text{value}}, \quad (6)$$

where $W_{\omega}^{\text{table}} \in \mathbb{R}^{1 \times K}$ is the attention values for ω and $W^{\text{value}} \in \mathbb{R}^{|\mathcal{X}| \times v}$ denotes the value weight with the value size $v \in \mathbb{N}$.

We note that a top- k operation is then performed such that only k mechanisms components are selected from the available K mechanisms (not selected mechanisms are zero masked), which enables an option to operate under an exponentially reduced belief state space by operating over only a subset of mechanisms.

Recurrent Layer Module. The recurrent layer module updates the hidden representation of all active mechanisms. Here, we have a separate RNN for each of the K mechanisms with all RNNs shared across all options. Each recurrent operation is done for each active mechanism separately by taking input $h_{\omega}^{\text{attn}} \in \mathbb{R}^{K \times v}$ and the previous RNN hidden state h_{ω}^{RNN} , in which we obtain the output of recurrent layer module $h_{\omega'}^{\text{RNN}} \in \mathbb{R}^{K \times h}$ for all mechanisms, where $h \in \mathbb{N}$ denotes RNN's hidden size.

Sparse Communication Module. The sparse communication module enables mechanisms to share information with one another, facilitating coordination in learning

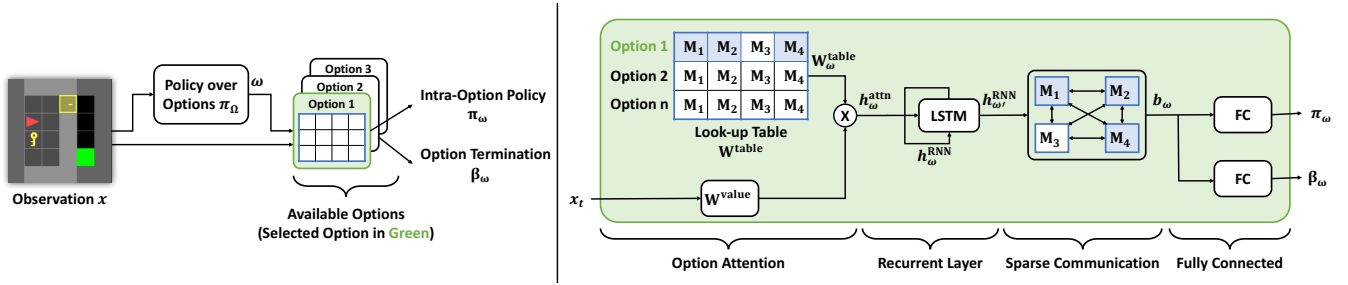


Figure 2: Architecture for intra-option policy of Context-Specific Abstracted Option-Critic (CRADOL). The left side describes the overview of CRADOL, and the right side details the option learning process.

amongst the options. Specifically, all active mechanisms can read hidden representations from both active and in-active mechanisms, with only active mechanisms updating their hidden states. This module outputs the context-specific belief state $b_\omega \in \mathbb{R}^{K \times h}$ given the input $h_{\omega'}^{\text{RNN}}$:

$$b_\omega = \left[\text{softmax} \left(\frac{W^{Q_{\text{comm}}} h_{\omega'}^{\text{RNN}} (W^{K_{\text{comm}}} h_{\omega'}^{\text{RNN}})^T}{\sqrt{d^{K_{\text{comm}}}}} \right) \right. \\ \left. W^{V_{\text{comm}}} h_{\omega'}^{\text{RNN}} \right] + h_{\omega'}^{\text{RNN}}, \quad (7)$$

where $W^{K_{\text{comm}}}$, $W^{Q_{\text{comm}}}$, $W^{V_{\text{comm}}}$ are communication parameters with the communication key size $d^{K_{\text{comm}}} \in \mathbb{N}$.

Equation (7) is similar to Goyal et al. (2021). We only update the top- k selected mechanisms from the option selection module, but through this transfer of information, each active mechanism can update its internal belief state and have improved exploration through contextualization. We note that there are many choices in what components to share across options. We explore these choices and their implications in our empirical analysis.

Fully Connected Module. Lastly, we have a separate fully connected layer for each option-specific intra-option π_ω and termination policy β_ω , which take b_ω as input. The intra-option policy provides the final output action a , and the termination determines whether ω should terminate.

Implementation

We first describe our choice for representing factored belief states. Under the state uniformity assumption discussed in Majeed and Hutter (2018), we assume the optimal policy network based on the agent’s history is equivalent to the network conditioned on the true state space. Hence, we refer to the representation learned as an implicit belief state. More explicit methods for modeling belief states have been considered, for example, as outlined in Igl et al. (2018). While the CRADOL framework is certainly compatible with this kind of explicit belief state modeling, we have chosen to implement the belief state implicitly and denote it as the factored belief state in order to have a fair empirical comparison to the RIMs method (Goyal et al. 2021) that we build off.

Our implementation draws inspiration from soft-actor critic (Haarnoja et al. 2018) to enable sample efficient off-policy optimization of the option-value functions, intra-option policy, and beta policy. In the appendix, we describe

the implementation of our approach (see ??) and additional details about the architecture.

Related Work

Hierarchical RL. There have been two major high-level approaches in the recent literature for achieving efficient HRL: the options framework to learn abstract skills (Sutton, Precup, and Singh 1999; Bacon, Harb, and Precup 2017; Omidshafiei et al. 2018; Riemer, Liu, and Tesauro 2018; Riemer et al. 2020) and goal-conditioned hierarchies to learn abstract sub-goals (Nachum et al. 2019; Levy, Platt, and Saenko 2019; Kulkarni and Narasimhan 2016; Kim et al. 2020). Goal-conditioned HRL approaches require a pre-defined representation of the environment and mapping of observation space to goal space, whereas the OC framework facilitates long-timescale credit assignment by dividing the problem into pieces and learning higher-level skills. Our work is most similar to Khetarpal et al. (2020) which learns a smaller initiation set through an attention mechanism to better focus an option to a region of the state space, and hence achieve specialization of skills. However, it does not leverage the focus to abstract and minimize the size of the belief space as CRADOL does. We consider a more explicit method of specialization by leveraging context-specific independence for representation abstraction. Our approach could also potentially consider learning initiation sets, so we consider our contributions to be orthogonal.

While OC provides a temporal decomposition of the problem, other approaches such as Feudal Networks (Vezhnevets et al. 2017) decompose problems with respect to the state space. Feudal approaches use a manager to learn more abstract goals over a longer temporal scale and worker modules to perform more primitive actions with fine temporal resolution. Our approach employs a combination of both visions, necessitating both temporal and state abstraction for effective decomposition. Although some approaches such as the MAXQ framework (Dietterich 2000) employ both, they involve learning recursively optimal policies that can be highly sub-optimal (Flet-Berliac 2019).

State Abstractions. Our approach is related to prior work that considers the importance of state abstraction and the decomposition of the learning problem (Konidaris 2019; Konidaris and Barto 2009; Jong and Stone 2005). Notable

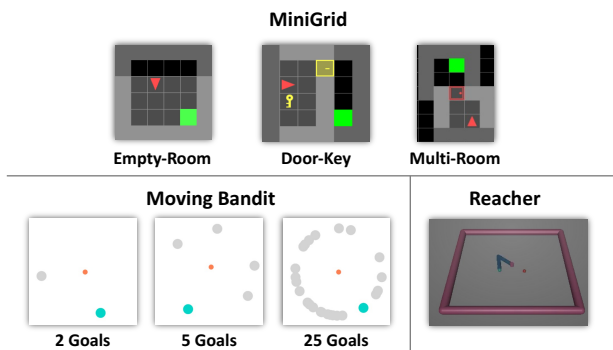


Figure 3: Visualization of MiniGrid (Chevalier-Boisvert, Willems, and Pal 2018), Moving Bandit (Frans et al. 2017), and Reacher (Brockman et al. 2016) domains.

methods of state abstraction include PAC state abstraction, which achieves correct clustering with high probability with respect to a distribution over learning problems (Abel et al. 2018). This abstraction method can have limited applicability to deep RL methods such as CRADOL. Zhang et al. (2019) has been able to learn task agnostic state abstractions by identifying casual states in the POMDP setting, whereas our approach considers discovering abstractions for sub-task specific learning. Konidaris, Kaelbling, and Lozano-Perez (2018) introduces the importance of abstraction in planning problems with Chitnis et al. (2020) performing a context-specific abstraction for the purposes of decomposition in planning-related tasks. CRADOL extends this work by exploring context-specific abstraction in HRL.

Evaluation

We demonstrate CRADOL’s efficacy on a diverse suite of domains. The code is available at <https://git.io/JucVH> and the videos are available at <https://bit.ly/3tpJc8Z>. We explain further details on experimental settings, including domains and hyperparameters, in the appendix.

Experimental Setup

Domains. We demonstrate the performance of our approach with domains shown in Figure 3. MiniGrid domains are well suited for hierarchical learning due to the natural emergence of skills needed to accomplish the goal. Moving Bandits considers the performance of CRADOL with extraneous features in sparse reward settings. Lastly, the Reacher domain observes the effects of CRADOL on low-level observation representations.

- **MiniGrid** (Chevalier-Boisvert, Willems, and Pal 2018): A library of open-source grid-world domains in sparse reward settings with image observations. Each grid contains exactly zero or one object with possible object types such as the wall, door, key, ball, box, and goal indicated by different colors. The goal for the domain can vary from obtaining a key to matching similar colored objects. The agent receives a sparse reward of 1 when it successfully reaches the green goal tile, and 0 for failure.

- **Moving Bandit** (Frans et al. 2017): This 2D sparse reward setting considers a number of marked positions in the environment that change at random at every episode, with 1 of the positions being the correct goal. An agent receives a reward of 1 and terminates when the agent is close to the correct goal position, and receives 0 otherwise.
- **Reacher** (Brockman et al. 2016): In this simulated MuJoCo task of OpenAI Gym environment, a robot arm consisting of 2 linkages with equal length must reach a random red target placed randomly at the beginning of each episode. We modify the domain to be a sparse reward setting: the agent receives a reward signal of 1 when its euclidean distance to the target is within a threshold, and 0 otherwise.

Baselines. We compare CRADOL to the following non-hierarchical, hierarchical, and modular recurrent neural network baselines:

- **A2C** (Mnih et al. 2016): This on-policy method considers neither context-specific nor temporal abstraction.
- **SAC** (Haarnoja et al. 2018): This entropy-maximization off-policy method considers neither context-specific nor temporal abstraction.
- **OC** (Bacon, Harb, and Precup 2017): We consider an off-policy implementation of OC based on the SAC method to demonstrate the performance of a hierarchical method considering only temporal abstraction.
- **A2C-RIM** (Goyal et al. 2021). This method considers A2C with recurrent independent mechanisms, a baseline that allows us to observe the performance of a method employing context-specific abstraction only.

Results

Question 1. *Does context-specific abstraction help achieve sample efficiency?*

To answer this question, we compare performance in the MiniGrid domains. With a sparse reward function and dense belief state representation (i.e., image observation), MiniGrid provides the opportunity to test the temporal and context-specific abstraction capabilities of our method. ?? shows both final task-level performance (\bar{V}) (i.e., final episodic average reward measured at the end of learning) and area under the learning curve (AUC). Higher values indicate better results for both metrics.

Overall, for all three scenarios of MiniGrid, CRADOL has the highest \bar{V} and AUC than the baselines. We observe that OC has a lower performance than CRADOL due to the inability of the options learned to diversify by considering the entire belief state space and the high termination probability of each option. Both A2C and SAC result in sub-optimal performance due to their failure in sparse reward settings. Finally, due to inefficient exploration and large training time required for A2C-RIM to converge, it is unable to match the performance of CRADOL. We see a smaller difference between CRADOL and these baselines for the Empty domain, as there is a smaller amount of context-specific abstraction that is required in this simplest setting

Algorithm	Abstraction		MiniGrid Empty		MiniGrid MultiRoom		MiniGrid KeyDoor	
	Temporal	State	\bar{V}	AUC	\bar{V}	AUC	\bar{V}	AUC
A2C	✗	✗	0.71±0.25	364±54	0.62±0.26	104±27	0.04±0.07	16±10
SAC	✗	✗	0.65±0.29	350±124	0.62±0.09	162±39	0.48±0.42	286±138
OC	✓	✗	0.56±0.36	312±143	0.40±0.15	108±45	0.55±0.44	252±217
A2C-RIM	✗	✓	0.57±0.40	283±47	0.19±0.23	52±22	0.57±0.40	12±1
CRADOL	✓	✓	0.92±0.12	470±21	0.75±0.06	200±27	0.87±0.17	499±25

Table 1: \bar{V} and Area under the Curve (AUC) in MiniGrid domains. Table shows mean and standard deviation computed with 10 random seeds. Best results in bold (computed by t -test with $p < 0.05$). Note that CRADOL has the highest \bar{V} and AUC compared to non-HRL (A2C, SAC), HRL (OC), and modular recurrent neural network (A2C-RIM) baselines. Figures 8-11 in the appendix show number of steps taken to converge for these experiments.

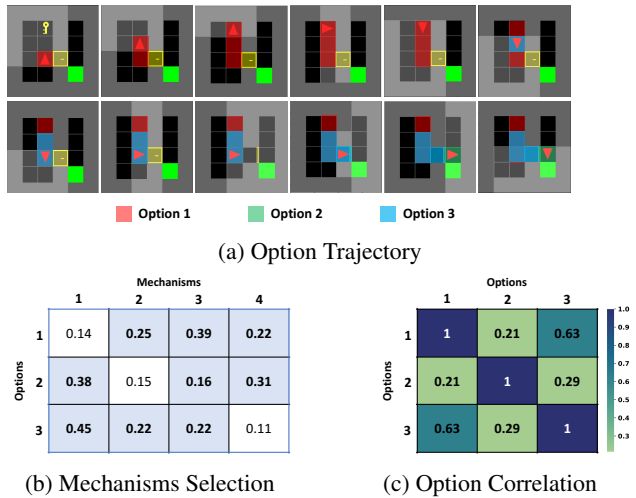


Figure 4: (a) Temporal abstraction by the use of 3 options for following tasks: option 1 for getting the key, option 2 for opening the door, and option 3 for going to the door. (b) Each option maps to a unique subset of mechanisms, corresponding to their unique functions in the domain. (c) There is low correlation between option 1 & 2 and option 2 & 3 but higher correlation between option 1 & 2 corresponding to the shared belief states between them.

of the MiniGrid domain. For the Multi-Room domain which is more difficult than the EmptyRoom domain, there is an increasingly larger gap as the agent needs to consider only the belief states in one room when trying to exit that room and belief states of the other room when trying to reach the green goal. Lastly, we see the most abstraction required for the Key Door domain where the baselines are unable to match the performance of CRADOL. As described in Figure 1, OC equipped with only temporal abstraction is unable to match the performance of CRADOL consisting of both temporal and context-specific abstraction.

Question 2. *What does it mean to have diverse options?*

We visualize the behaviors of options for temporal abstraction and mechanisms for context-specific abstraction to further understand whether options are able to learn diverse sub-policies. Figure 4a shows the option trajectory in the

DoorKey domain for the following (learned) sub-tasks: getting the key, opening the door, and going to the door. We find that each option is only activated for one sub-task. Figure 4b shows the mapping between options and mechanisms, and we see that each option maps to a unique subset of mechanisms. To understand whether these mechanisms have mapped to different factors of the belief state (and hence have diverse parameters), Figure 4c computes the correlation between options, measured by the Pearson product-moment correlation method (Freedman, Pisani, and Purves 2007). We find low correlation between option 1 & 2 and option 2 & 3 but higher correlation between option 1 & 3. Specifically, we observe a high correlation between option 1 & 3 in getting the key and opening the door due to the shared states between them, because opening the door is highly dependent on obtaining the key in the environment. This visualization empirically verifies our hypothesis that both temporal and context-specific abstraction are necessary to learn diverse and complementary option policies.

Question 3. *How does performance change as the need for context-specific abstraction increases?*

In order to understand the full benefits of context-specific abstraction, we observe the performance with increasing context-specific abstraction in the Moving Bandits domain. We consider the performance determined by AUC for an increasing number of spurious features in the observation. Namely, we add 3 & 23 additional goals to the original 2 goal observation to test the capabilities of CRADOL (see Figure 3). Increasing the number of goals requires an increasing amount of context-specific abstraction, as there are more spurious features the agent must ignore to learn to move to the correct goal location as indicated in its observation. As shown in Figure 5, CRADOL performs significantly better than OC as the number of spurious goal locations it must ignore increases. We expect that this result is due to the CRADOL’s capability of both temporal and context-specific abstraction.

Question 4. *What is shared between options in the context of mechanisms?*

As described in Figure 2, there are 4 components that make up the option learning model. In order to investigate which are necessary to share between options, we perform an ablation study with a subset of the available choices shown




Figure 5: AUC between CRADOL and OC in Moving Bandit domain. As the number of spurious features on the x-axis increases, the gap between the AUC performance between CRADOL and OC increases. This indicates a greater need for context-specific abstraction. Note that we see a different AUC scale across the number of goals simply due to different max train iteration for each goal setting. Mean and 95% confidence interval computed for 10 seeds are shown.

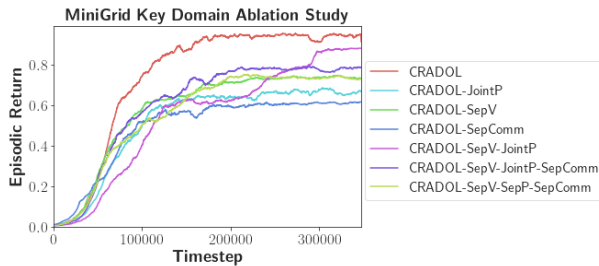


Figure 6: Ablation study experimenting with various components that can be shared and not shared between options in the option learning process shown in Figure 2. Too much sharing or too little sharing between options can lead to sub-optimal performance due to a lack of coordination. This result shows the mean computed for 10 seeds.

in Figure 6. Specifically, we study sharing of the look-up table W^{table} in the input attention module between options (CRADOL-JointP), learning a separate parameter W^{value} between options (CRADOL-SepV), learning separate parameters for ($W^{Q_{\text{comm}}}$, $W^{K_{\text{comm}}}$, $W^{V_{\text{comm}}}$) of the sparse communication layer for each option (CRADOL-SepComm), and three other combinations of these three modules.

We find the lowest performance for the method with a separate sparse communication module for each option. We hypothesize that this is due to a lack of coordination between each option in updating their active mechanisms and ensuring other non-active mechanisms are learning separate and complementary parameters. Having a joint look-up table W^{table} results in the second-lowest performance. This effectively maps each option to the same set of mechanisms, leading to a lack of diversity between option policies and only allowing for the difference in the fully connected layer of each option. Lastly, we observe the third-lowest performance with a separate parameter W^{value} between options. Each option learning from a different representation can lead




Figure 7: In domains with small or negligible context-specific abstraction, the benefit of CRADOL is not as significant compared to the Reacher domain, where the low-dimensional observation representation does not require context-specific abstraction. This figure shows the mean and 95% confidence interval computed for 10 seeds.

to similar option-policies and factored belief states for certain sub-goals unaccounted for. Other combinations reaffirm that sharing the sparse communication larger is essential for coordinated behavior when learning option policies.

Question 5. *Is context-specific abstraction always beneficial?*

The Reach domain allows us to observe the effects of our method when there is little benefit of reducing the problem size, namely, when the observation is incredibly small. This domain does not require context-specific abstraction as the entire belief state space consists of relevant information to achieve the goal at hand. Specifically, the low-level representation of the observation as a 4-element observation vector, with the first 2 elements containing the generalized positions and the next 2 elements containing the generalized velocities of the two arm joints, are essential to reach the goal location. As expected in Figure 7, the performance between CRADOL and OC is similar in this domain as the observation space does not contain any features that are useful for CRADOL to perform context-specific abstraction. We note our gains are larger for problems where they have larger intrinsic dimensionality.

Conclusion

In this paper, we have introduced Context-Specific Representation Abstraction for Deep Option Learning (CRADOL) for incorporating both temporal and state abstraction for the effective decomposition of a problem into simpler components. The key idea underlying our proposed formulation is to map each option to a reduced state space, effectively considering each option as a subset of mechanisms. We evaluated our method on several benchmarks with varying levels of required abstraction in sparse reward settings. Our results indicate that CRADOL is able to decompose the problems at hand more effectively than state-of-the-art HRL, non-HRL, and modular recurrent neural network baselines. We hope that our work can help provide the community with a theoretical foundation to build off for addressing the deficiencies in HRL methods.

Acknowledgements

Research funded by IBM, Samsung (as part of the MIT-IBM Watson AI Lab initiative) and computational support through Amazon Web Services.

References

- Abel, D.; Arumugam, D.; Lehnert, L.; and Littman, M. 2018. State Abstractions for Lifelong Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 10–19. PMLR.
- Bacon, P.-L.; Harb, J.; and Precup, D. 2017. The Option-Critic Architecture. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, AAAI’17, 1726–1734. AAAI Press.
- Boutillier, C.; Friedman, N.; Goldszmidt, M.; and Koller, D. 1996. Context-Specific Independence in Bayesian Networks. In *Proceedings of the Twelfth International Conference on Uncertainty in Artificial Intelligence*, UAI’96, 115–123. Morgan Kaufmann Publishers Inc.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. *CoRR*, abs/1606.01540.
- Chevalier-Boisvert, M.; Willems, L.; and Pal, S. 2018. Minimalistic Gridworld Environment for OpenAI Gym. <https://github.com/maximecb/gym-minigrid>.
- Chitnis, R.; Silver, T.; Kim, B.; Kaelbling, L. P.; and Lozano-Pérez, T. 2020. CAMPs: Learning Context-Specific Abstractions for Efficient Planning in Factored MDPs. *CoRR*, abs/2007.13202.
- Daniel, C.; Van Hoof, H.; Peters, J.; and Neumann, G. 2016. Probabilistic Inference for Determining Options in Reinforcement Learning. 104(2–3).
- Dietterich, T. 2000. Hierarchical Reinforcement Learning with the MAXQ Value Function Decomposition. *The Journal of Artificial Intelligence Research (JAIR)*, 13.
- Flet-Berliac, Y. 2019. The Promise of Hierarchical Reinforcement Learning. *The Gradient*.
- Frans, K.; Ho, J.; Chen, X.; Abbeel, P.; and Schulman, J. 2017. Meta Learning Shared Hierarchies. *CoRR*, abs/1710.09767.
- Freedman, D.; Pisani, R.; and Purves, R. 2007. Statistics (international student edition). *Pisani, R. Purves, 4th edn. WW Norton & Company, New York*.
- Goyal, A.; Lamb, A.; Hoffmann, J.; Sodhani, S.; Levine, S.; Bengio, Y.; and Schölkopf, B. 2021. Recurrent Independent Mechanisms. In *International Conference on Learning Representations*.
- Guestrin, C.; Koller, D.; Parr, R.; and Venkataraman, S. 2003. Efficient Solution Algorithms for Factored MDPs. *Journal of Artificial Intelligence Research*, 19: 399–468.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1861–1870. PMLR.
- Igl, M.; Zintgraf, L.; Le, T. A.; Wood, F.; and Whiteson, S. 2018. Deep Variational Reinforcement Learning for POMDPs. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 2117–2126. PMLR.
- Jong, N. K.; and Stone, P. 2005. State Abstraction Discovery from Irrelevant State Variables. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, IJCAI’05, 752–757. Morgan Kaufmann Publishers Inc.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2): 99–134.
- Kamat, A.; and Precup, D. 2020. Diversity-Enriched Option-Critic. *CoRR*, abs/2011.02565.
- Khetarpal, K.; Klissarov, M.; Chevalier-Boisvert, M.; Bacon, P.-L.; and Precup, D. 2020. Options of interest: Temporal abstraction with interest functions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 4444–4451.
- Kim, D.-K.; Liu, M.; Omidshafiei, S.; Lopez-Cot, S.; Riemer, M.; Habibi, G.; Tesauro, G.; Mourad, S.; Campbell, M.; and How, J. P. 2020. Learning Hierarchical Teaching Policies for Cooperative Agents. *AAMAS ’20*, 620–628. International Foundation for Autonomous Agents and Multiagent Systems.
- Konidaris, G. 2019. On the necessity of abstraction. *Current Opinion in Behavioral Sciences*, 29: 1–7. *Artificial Intelligence*.
- Konidaris, G.; and Barto, A. 2009. Efficient Skill Learning Using Abstraction Selection. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, IJCAI’09, 1107–1112. Morgan Kaufmann Publishers Inc.
- Konidaris, G.; Kaelbling, L. P.; and Lozano-Perez, T. 2018. From Skills to Symbols: Learning Symbolic Representations for Abstract High-Level Planning. *J. Artif. Int. Res.*, 215–289.
- Kulkarni, T. D.; and Narasimhan, K. R. 2016. Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation. *Neural Information Processing Systems 2016*.
- Levy, A.; Platt, R.; and Saenko, K. 2019. Hierarchical Reinforcement Learning with Hindsight. In *International Conference on Learning Representations*.
- Majeed, S. J.; and Hutter, M. 2018. On Q-learning Convergence for Non-Markov Decision Processes. 2546–2552.
- Mankowitz, D. J.; Mann, T. A.; and Mannor, S. 2016. Adaptive Skills Adaptive Partitions (ASAP). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous Methods for Deep Reinforcement Learning. In

Proceedings of The 33rd International Conference on Machine Learning, volume 48 of *Proceedings of Machine Learning Research*, 1928–1937. PMLR.

Nachum, O.; Gu, S.; Lee, H.; and Levine, S. 2019. Near-Optimal Representation Learning for Hierarchical Reinforcement Learning. In *International Conference on Learning Representations*.

Omidshafiei, S.; Kim, D.-K.; Pazis, J.; and How, J. P. 2018. Crossmodal Attentive Skill Learner. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '18, 139–146. International Foundation for Autonomous Agents and Multiagent Systems.

Precup, D.; and Sutton, R. S. 2000. *Temporal Abstraction in Reinforcement Learning*. Ph.D. thesis.

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. USA: John Wiley & Sons, Inc., 1st edition. ISBN 0471619779.

Riemer, M.; Cases, I.; Rosenbaum, C.; Liu, M.; and Tesauro, G. 2020. On the role of weight sharing during deep option learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 5519–5526.

Riemer, M.; Liu, M.; and Tesauro, G. 2018. Learning Abstract Options. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.

Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1): 181–211.

Vezhnevets, A.; Mnih, V.; Osindero, S.; Graves, A.; Vinyals, O.; Agapiou, J.; and kavukcuoglu, k. 2016. Strategic Attentive Writer for Learning Macro-Actions. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.

Vezhnevets, A. S.; Osindero, S.; Schaul, T.; Heess, N.; Jaderberg, M.; Silver, D.; and Kavukcuoglu, K. 2017. FeUdal Networks for Hierarchical Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 3540–3549. PMLR.

Zhang, A.; Lipton, Z. C.; Pineda, L.; Azizzadenesheli, K.; Anandkumar, A.; Itti, L.; Pineau, J.; and Furlanello, T. 2019. Learning Causal State Representations of Partially Observable Environments. *CoRR*, abs/1906.10437.