

Hierarchical Reinforcement Learning: A Comprehensive Survey

SHUBHAM PATERIA, Nanyang Technological University

BUDHITAMA SUBAGDJA, Singapore Management University

AH-HWEE TAN, Singapore Management University

CHAI QUEK, Nanyang Technological University

Hierarchical Reinforcement Learning (HRL) enables autonomous decomposition of challenging long-horizon decision-making tasks into simpler subtasks. During the past years, the landscape of HRL research has grown profoundly, resulting in copious approaches. A comprehensive overview of this vast landscape is necessary to study HRL in an organized manner. We provide a survey of the diverse HRL approaches concerning the challenges of learning hierarchical policies, subtask discovery, transfer learning, and multi-agent learning using HRL. The survey is presented according to a novel taxonomy of the approaches. Based on the survey, a set of important open problems are proposed to motivate the future research in HRL. Furthermore, we outline a few suitable task domains for evaluating the HRL approaches and a few interesting examples of the practical applications of HRL, in the Supplementary Material.

CCS Concepts: • **Computing methodologies** → **Reinforcement learning**.

Additional Key Words and Phrases: hierarchical reinforcement learning, subtask discovery, skill discovery, hierarchical reinforcement learning survey, hierarchical reinforcement learning taxonomy

ACM Reference Format:

Shubham Pateria, Budhitama Subagdja, Ah-Hwee Tan, and Chai Quek. 2023. Hierarchical Reinforcement Learning: A Comprehensive Survey. 1, 1 (August 2023), 35 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

One of the cardinal goals of Artificial Intelligence is to develop autonomous agents which can perform various complex tasks in an environment by planning the optimal sequences of actions [31]. Reinforcement Learning (RL) is a computational paradigm for *learning* a policy that takes optimal actions in various states of a task environment, so as to maximize the cumulative rewards received by the acting agent [83]. In order to learn the optimal policy, the agent explores the state and action spaces relevant to the task by executing various sequences of *state-action-next state* transitions. The average length of such sequences¹ is called the task horizon. If the horizon is long while the task involves large state and action spaces, then the exploration space also becomes large.

¹The length of a sequence is defined as the number of items in that sequence. In this case, an item is an action.

Authors' addresses: Shubham Pateria, SHUBHAM007@e.ntu.edu.sg, School of Computer Science and Engineering, Nanyang Technological University, 50 Nanyang Avenue, Singapore, 639798; Budhitama Subagdja, budhitamas@smu.edu.sg, School of Computing and Information Systems, Singapore Management University, 80 Stamford Road, Singapore, 178902; Ah-Hwee Tan, ahtan@smu.edu.sg, School of Computing and Information Systems, Singapore Management University, 80 Stamford Road, Singapore, 178902; Chai Quek, ashcquek@ntu.edu.sg, School of Computer Science and Engineering, Nanyang Technological University, 50 Nanyang Avenue, Singapore, 639798.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

XXXX-XXXX/2023/8-ART \$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

This results in the poor performance of the standard RL algorithms [4, 41, 52] on such *long-horizon tasks* without sophisticated exploration techniques [93, 44, 64].

Hierarchical Reinforcement Learning (HRL) decomposes a long-horizon reinforcement learning task into a hierarchy of subproblems or subtasks such that a higher-level policy learns to perform the task by choosing optimal subtasks as the higher-level actions. A subtask may itself be a reinforcement learning problem with a lower-level policy learning to solve it [34]. This hierarchy of policies collectively determines the behaviour of the agent. Task decomposition effectively reduces the original task's long horizon into a shorter horizon in terms of the sequences of subtasks. This is because each subtask is a higher-level action which persists for a longer timescale compared to a lower-level action, a property which is often referred to as temporal abstraction [11, 20, 12]. Temporal abstraction can also enable efficient credit assignment over longer timescales [67]. At the same time, a subtask may itself be easier to learn and the learned subtasks lead to more structured exploration over the course of training of the HRL agent [93].

These aspects make HRL a promising approach to scale reinforcement learning to the long-horizon tasks [5, 12, 67, 76]. HRL algorithms have been shown to outperform standard RL in several long-horizon problems such as continuous control² [91, 76, 57], long-horizon games [50, 67], robot manipulation [59, 99] etc. Different empirical studies find that the performance benefits of HRL are primarily due to improved exploration using subtasks/subgoals [28, 93].

HRL research has progressed significantly in the past three decades, resulting in a profusion of approaches that address a diversity of challenges such as learning the policies in a hierarchy, autonomous discovery of subtasks, transfer learning, and multi-agent learning using HRL. This causes significant difficulties in understanding the progress in the field in an organized manner. Hence, there is a necessity for a comprehensive survey to collect and organize the important HRL approaches, and to provide a general taxonomy for their classification.

What are the differences between this survey and the previous ones? Barto et al [20] provided a survey of the advances in HRL up to the year 2003. That survey includes an important overview of the classical approaches, mainly MAXQ [12], Options [11], and HAMs [10]. Significant developments have occurred in the field of HRL since the time of that survey, such as subtask discovery using graph analysis, variational inference, autoencoding, unified HRL, subtask discovery in multi-agent HRL, transfer learning with HRL etc. Our survey mainly differs from that of Barto et al [20] in the sense that we review the new HRL approaches which have emerged after their survey in addition to the classical approaches covered by them. Al-Emran et al [40] performed a survey from the perspective of the practical applications of HRL. However, it does not include several important approaches which do not apply to the chosen applications, such as the recent unified HRL techniques, transfer learning with HRL, multi-agent HRL etc. In contrast, we survey the HRL approaches broadly and provide a general taxonomy which is application-agnostic.

The most recent survey by Mendonca et al [92] gives a detailed review of the graph-based approaches for subtask discovery. Due to this limited scope, their survey goes into the depths of the graph-based subtask discovery but excludes other important aspects of HRL research such as the techniques for learning a hierarchy of policies, subtask discovery using variational inference, unified HRL, transfer learning with HRL, multi-agent HRL etc. On the other hand, we review all such approaches along with the graph-based subtask discovery (subsection 3.3.1) in the spirit of providing a broader view of the HRL research.

²An agent uses a continuous action space instead of a discrete set of actions.

Contributions of this survey. The objective of this survey is to thoroughly review the literature on Hierarchical Reinforcement Learning (HRL) and provide a panorama of the approaches developed so far. The key contributions are as follows:

- (1) We conduct a comprehensive survey of the work done so far in the field of HRL. The survey includes the approaches for learning hierarchies of policies, independent subtask discovery, unified HRL, multi-task/transfer learning with HRL, and multi-agent HRL.
- (2) We provide a novel taxonomy to organize the HRL approaches along important characteristic dimensions such as single-agent vs multi-agent, single-task vs multi-task, and without subtask discovery vs with subtask discovery.
- (3) We identify a set of important open problems to provide the directions for future research concerning the scalability, efficiency, and theoretical robustness of HRL.

Our endeavour is to make this survey as comprehensive as possible with regard to the overview and classification of the HRL research in general, rather than the review of all the specific approaches in existence. The HRL literature is too vast to cover in a limited number of pages. Therefore, the main representative approaches for each class have been included in the survey while other similar approaches have been left out to conserve space.

The rest of the paper is organized as follows. In section 2, we review the general concepts of reinforcement learning, task decomposition, and HRL. Section 3 presents the taxonomy and survey of various HRL approaches along with the broad challenges addressed by them. Section 4 contains a discussion on the important open problems for future research on HRL. Finally, we conclude the survey in section 5. In addition to this paper, we provide a Supplementary Material document which outlines a few task domains for evaluating and benchmarking HRL approaches, and contains a few interesting examples of the practical applications of HRL in the real-world problems.

2 PRELIMINARIES

2.1 Reinforcement Learning

Reinforcement Learning (RL) refers to learning how to *take actions in different states of an environment* so as to *maximize cumulative future reward* [83], where the reward is a form of feedback associated with a task and received in response to an action. The majority of the current RL research is built upon the theory of Markov Decision Processes (MDPs) [1]. An MDP is control process defined using a tuple $\langle S, A, P, r \rangle$. Here, S is the state space of the environment, where a state $s \in S$ represents a situation of the environment. A is the set of actions which can be taken by an agent. $r : S \times A \rightarrow \mathbb{R}$ is the numerical reward obtained as a function of state and action. The effect of an action on the future state is captured by a probabilistic transition function $P(s_{t+1}|s_t, a_t) \rightarrow [0, 1]$ where $s_{t+1} \in S$ is the next state observed after taking action $a_t \in A$ in a state $s_t \in S$ at time t . An MDP is stochastic if there exists a state-action pair (s_t, a_t) for which $P(s_{t+1}|s_t, a_t) \neq 1$, because the effect of an action is not deterministic. Contrarily, an MDP is deterministic if for every state-action pair (s_t, a_t) , $P(s_{t+1}|s_t, a_t) = 1$. The agent takes actions according to a policy $\pi : S \times A \rightarrow [0, 1]$. This policy may be stochastic.

One step of interaction between the agent and the environment produces a transition from current state s_t to next state s_{t+1} , observed as an experience data sample (s_t, a_t, r_t, s_{t+1}) . A sequence of such transitions constitutes a trajectory. The trajectory taken from any starting state and action pair (s_t, a_t) is affected by the stochasticity in the policy and in the state transitions (as per $P(s_{t+1}|s_t, a_t)$). Thus, the value of taking a_t in s_t is formally calculated as the *expected* (or average) cumulative

reward obtained over all possible trajectories, as follows:

$$Q^\pi(s_t, a_t) = r(s_t, a_t) + \mathbb{E}_{a \sim \pi(s)} \left[\sum_{i=1}^{\infty} \gamma^i r(s_{t+i}, a_{t+i}) | s_t, a_t \right] \quad (1)$$

Here, $Q^\pi(s_t, a_t)$ is called the *Q-value* of the state-action pair while an agent follows a policy π . $\gamma \in [0, 1)$ is called the *discount factor*. $a \sim \pi(s)$ implies that an action is sampled using the policy π . **The core objective of RL** is to learn an optimal policy π^* which satisfies the following condition:

$$\pi^* = \arg \max_{\pi} Q^\pi(s, a), \forall s \in S, \forall a \in A \quad (2)$$

2.2 Hierarchical Reinforcement Learning

It can be inferred from equations 1 and 2 that the objective of an RL agent is to search for a policy which maximizes the cumulative reward averaged over various possible trajectories the agent can take while following that policy. While exploring the state and action spaces to learn the optimal policy, the agent takes various trajectories whose expected length is the task horizon (the horizon is shown as theoretically infinite in equation 1, but in practice it may be finite). As mentioned in section 1, when the state and action spaces are large, and the task horizon is long, the exploration becomes challenging using the standard RL approach [93, 44, 64].

Hierarchical Reinforcement Learning (HRL) provides a mechanism to perform a challenging task by decomposing it into simpler subtasks using a hierarchy of policies learned via reinforcement learning. In such a hierarchy, the highest-level policy generally chooses the subtasks of the main task as its actions [34, 12]. This policy is trained to perform the main task in terms of the sequence of its subtasks, using the rewards obtained in the main task. At a lower-level in the hierarchy, a subtask chosen by the higher-level is itself a reinforcement learning problem. A lower-level policy learns to perform that subtask using the internal rewards related to it (optionally, the main task reward may also be added). The lowest-level policies choose the basic actions which are henceforth referred to as *primitive* actions.

The process of HRL is illustrated in Fig. 1 using an example of task decomposition given by Russell et al [31]. In this illustration, an HRL agent decomposes and performs a long-horizon task "Going to Hawaii" (GTH). The HRL agent is composed of a hierarchy of policies. The task policy π_{GTH} decomposes the original task GTH into a sequence of highest-level subtasks "Book Tickets" (BT), "Go To Airport" (GTA), and so on. The task policy initially chooses BT. Then, BT is executed for multiple time steps until its termination occurs at time T_3 . During this period, the policy of the subtask BT, i.e. π_{BT} , itself chooses different shorter subtasks sequentially. These are "Open Booking Website" (OBW), "Enter Flight Information" (EFI), and so on. After BT terminates at T_3 , the task policy chooses GTA, which itself chooses a shorter subtask "Go to Taxi Stand" (GTS). During each time step, a primitive action a is chosen by the lowest level subtask policy, for e.g. by π_{OBW} , π_{EFI} , π_{GTS} etc. The HRL agent receives a task reward r^{GTH} in response to the primitive action, which is accumulated and given at different time-scales to different levels. The task reward may be optional for the policies below the highest level, which can be trained using the subtask-related rewards. In this manner, the HRL process continues in time until GTH finishes.

2.2.1 Formal definition of a subtask. The definition of a subtask in the formal context of HRL is provided in this subsection. This definition is a culmination of the different ways in which different HRL approaches define the subtasks. Hence, it should be considered as a *general* definition. A particular HRL approach may only use selective components of the general definition. The specifics are further discussed in section 3.

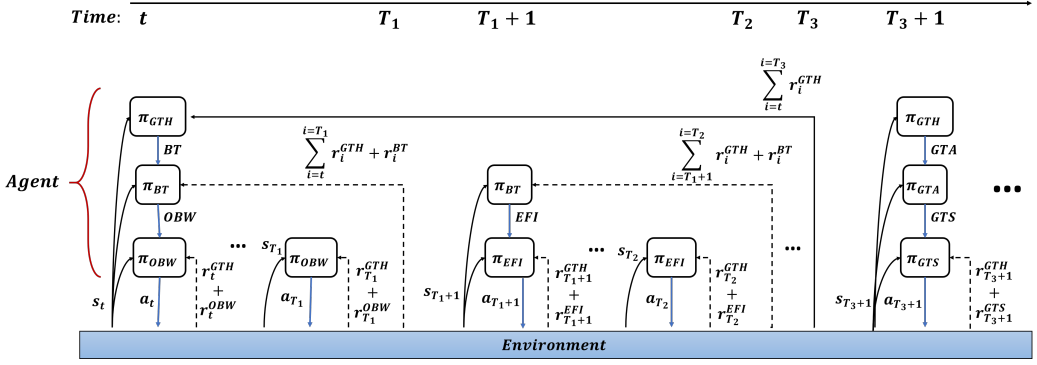


Fig. 1. The temporal process of Hierarchical Reinforcement Learning.

Firstly, we denote the main long-horizon task as Γ and the task policy as π_Γ . The task policy is at the top of a hierarchy, for e.g. π_{GTH} in Fig. 1. A subtask is denoted as ω . It is defined using the components described as follows.

- π_ω , which is the **policy of the subtask**. It maps the environment states to primitive actions or to the subtasks of ω [11, 12].
- The **objective components**:
 - r_ω , which is the subtask reward used to train π_ω . This is typically different from the reward associated with the main task [12, 5, 76, 67, 50],
 - g_ω , which is a subgoal or a set of subgoals associated with ω . A subgoal might be a state $s \in S$ itself [15, 29, 19], an abstract form of a state [76, 91], a learned embedding [67, 77, 82] etc. The reward r_ω might be defined with respect to the subgoal(s).
- The **execution components**:
 - I_ω , which is the initiation condition of ω . It may be defined as a set of states in which ω can be chosen for execution [11, 12], as a function which modulates the probability of choosing ω in a particular state [101], as a set of logical conditions [94] etc.
 - β_ω , which is the termination condition of ω . It may be defined as a set of states in which ω should terminate if it is being executed [12, 5, 50], as a function which modulates the probability of terminating ω in a particular state [11, 56], as a fixed time-limit [91, 76] etc. If g_ω is defined, then the subgoal state is also usually included as a state in which ω must terminate [12, 50].

In the remainder of this paper, the term 'subtask' should be inferred as referring to this formal definition.

2.2.2 Formalism of HRL based on Semi-markov Decision Process. HRL is formalized on the basis of the theory of Semi-Markov Decision Process (SMDP) [32]. An SMDP is a stochastic control process similar to an MDP (subsection 2.1), but unlike MDP, it also involves the concept of *time* for which an action is executed after it has been chosen. In the context of HRL, the actions with the concept of time are the *subtasks*. Starting from a state $s_t \in S$, assume that an agent chooses a subtask $\omega_t \in \Omega$, where Ω is the set of subtasks (or the subtask space). Then, the transition function of the SMDP is defined as a joint distribution

$$P(s_{t+c_{\omega_t}}, c_{\omega_t} | s_t, \omega_t) = P(s_{t+c_{\omega_t}} | s_t, \omega_t, c_{\omega_t}) P(c_{\omega_t} | s_t, \omega_t) \quad (3)$$

Here, c_{ω_t} denotes the number of time steps for which ω_t is executed, starting from the state s_t . c_{ω_t} is actually determined by the termination condition β_{ω_t} which is one of the execution components defined in subsection 2.2.1.

The reward obtained in response to performing the subtask ω_t starting from state s_t is denoted as $R(s_t, \omega_t)$, calculated as follows,

$$R(s_t, \omega_t) = \mathbb{E}_{a \sim \pi_{\omega_t}(s)} \left[\sum_{i=0}^{c_{\omega_t}-1} \gamma^i r(s_{t+i}, a_{t+i}) | s_t, a_t = \pi_{\omega_t}(s_t) \right] \quad (4)$$

Equation 4 indicates that the reward $R(s_t, \omega_t)$ is equal to the expected cumulative reward obtained while following the subtask policy π_{ω_t} from time t until the termination of ω_t after c_{ω_t} time steps. Now, an optimal task policy would be the one which leads to the following desired maximum Q-value:

$$Q(s_t, \omega_t) = R(s_t, \omega_t) + \sum_{s_{t+c_{\omega_t}}, c_{\omega_t}} \gamma^{c_{\omega_t}} P(s_{t+c_{\omega_t}}, c_{\omega_t} | s_t, \omega_t) \max_{\omega_{t+c_{\omega_t}}} Q(s_{t+c_{\omega_t}}, \omega_{t+c_{\omega_t}}) \quad (5)$$

$\forall s \in S$ and $\forall \omega \in \Omega$.

It should be noted that the Q-value in equation 5 also depends on $R(s_t, \omega_t)$ and $P(s_{t+c_{\omega_t}}, c_{\omega_t} | s_t, \omega_t)$. These two quantities are determined by the execution of ω_t using its policy π_{ω_t} . Therefore, an agent actually needs to learn multiple policies at different levels of a task decomposition hierarchy, including π_{Γ} and the policies of all the subtasks (e.g., Fig. 1). We extend the notations of subtasks and policies for a multi-level hierarchy as defined below. Please refer to Fig. 2 for an example of these notations in the context of an HRL agent with a three level hierarchy.

- ω^l : a subtask at level l of the hierarchy,
- Ω_{ω^l} : set of subtasks under the subtask ω^l such that $\omega^{l-1} \in \Omega_{\omega^l}$,
- $\pi_{\omega^l}: S \times \Omega_{\omega^l} \rightarrow [0, 1]$: policy of the subtask ω^l . In other words, ω^{l-1} is chosen by π_{ω^l} ,
- $\Omega_{\omega^1} = A$, i.e., the output space of a subtask at the lowest level ($l = 1$) is the primitive action space A ,
- π_{Γ} and Ω_{Γ} denote the main task policy and the set of subtasks at the highest-level, respectively.

Consolidating all the definitions provided above, there are two principal parts of an HRL agent:

- **Subtask space** ($\Omega_{hierarchy}$). This is the super-set of all the subtasks used in a hierarchy, i.e. $\Omega_{hierarchy} = \{\Omega_{\omega^2}, \Omega_{\omega^3}, \Omega_{\omega^4}, \dots, \Omega_{\Gamma}\}$.
- **Hierarchical policy** ($\pi_{hierarchy}$). The primitive action taken by the HRL agent is the result of the recursive choices of subtasks. Consider the three level hierarchy depicted in Fig. 2. In this hierarchy, the main policy π_{Γ} chooses a level 2 subtask, i.e. $\omega^2 = \pi_{\Gamma}(s)$ where $\omega^2 \in \Omega_{\Gamma}$. The policy of ω^2 is executed until its termination as per β_{ω^2} . It chooses the lowest-level subtask $\omega^1 = \pi_{\omega^2}(s)$ where $\omega^1 \in \Omega_{\omega^2}$. The policy of ω^1 is executed until its termination as per β_{ω^1} . This lowest-level policy selects a primitive action, i.e. $a = \pi_{\omega^1}(s)$. This complete *state-to-subtask-to-action mapping* from π_{ω^1} is called the *hierarchical policy*, denoted as $\pi_{hierarchy}$. Now, the primitive action taken by the HRL agent *as a whole* can be equivalently expressed as $a = \pi_{hierarchy}(s)$. This description can be extrapolated to the hierarchies with more than three levels.

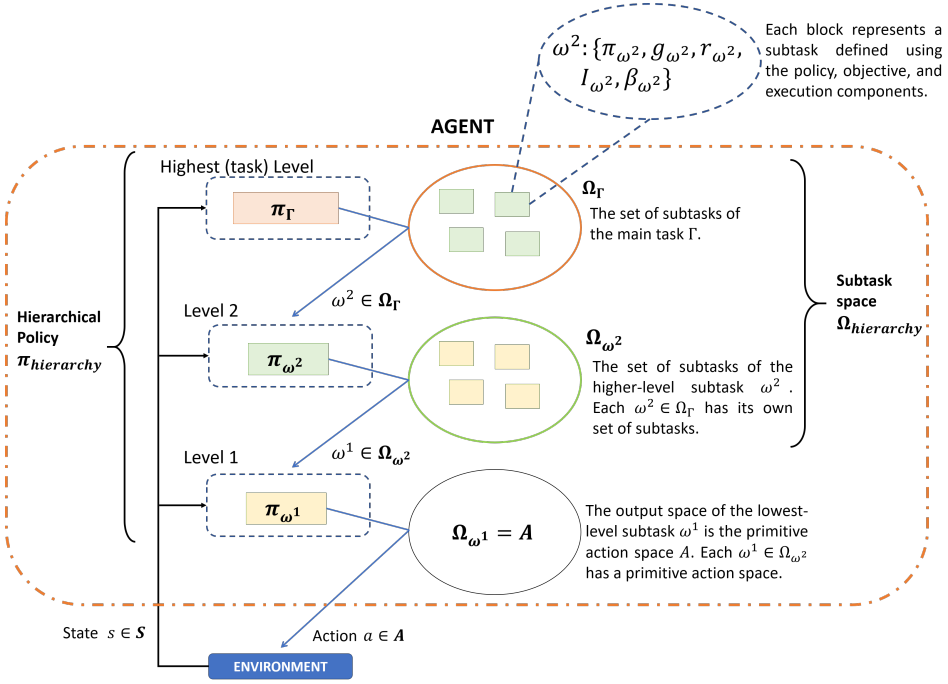


Fig. 2. **The concept of a Hierarchical Reinforcement Learning agent.** This illustration is for a three level hierarchy, but the concept can be extrapolated to more than three levels.

Based on the definitions provided above, the expected discounted cumulative reward received by the HRL agent can be written as,

$$Q^{hierarchy}(s_t, a_t) = \mathbb{E}_{a \sim \pi_{hierarchy} | \Omega_{hierarchy}} \left[\sum_{i=0}^{\infty} \gamma^{t+i} r(s_{t+i}, a_{t+i}) | s_t, a_t \right] \quad (6)$$

where $a \sim \pi_{hierarchy} | \Omega_{hierarchy}$ indicates that a primitive action a is sampled using the hierarchical policy $\pi_{hierarchy}$ conditioned on the available subtask space $\Omega_{hierarchy}$.

2.2.3 Problem definition of HRL. The general problem definition of HRL is to find the optimal hierarchical policy $\pi_{hierarchy}^*$ and the optimal subtask space $\Omega_{hierarchy}^*$ as the solution to:

$$\Omega_{hierarchy}^*, \pi_{hierarchy}^* = \underset{\Omega_{hierarchy}}{\operatorname{argmax}} \underset{\pi_{hierarchy} | \Omega_{hierarchy}}{\operatorname{argmax}} Q^{hierarchy}(s, a), \forall s \in S, a \in A \quad (7)$$

The HRL problem expressed in equation 7 can be divided into two parts. The first part is **learning hierarchical policy** which refers to finding the optimal hierarchical policy conditioned on the available subtask space (i.e. $\underset{\pi_{hierarchy} | \Omega_{hierarchy}}{\operatorname{argmax}}$). This is essential because the hierarchical policy determines the behaviour of an HRL agent on any given task. The policies at various levels of $\pi_{hierarchy}$ can be learned simultaneously in an end-to-end manner [91, 76, 51, 12, 56] or they may be learned one level at a time, in a bottom-to-top manner [59, 63, 57]. The second part is **subtask discovery** which refers to automatically finding the optimal subtask space using the HRL agent's experience data (i.e. $\underset{\Omega_{hierarchy}}{\operatorname{argmax}}$). Subtask discovery is not essential because a subtask space could be handcrafted using precise domain knowledge. However, it is necessary for generalized HRL without dependence on the manual handcrafting.

2.3 Definitions of common terms and concepts

A few important terms and concepts which appear commonly in the rest of this paper are defined here.

Skill. A 'skill' is a semantic term for referring to a primitive action policy learnt to perform a subtask, in the sense that such a policy represents the *ability of an agent to do something well*.

Universal policy. In this paper, a 'universal policy' means a policy which can learn all the possible subtasks by taking the corresponding subgoals [5, 67, 91] or instructions [89] as input, i.e. $\pi_\omega(s) = \pi(s, g_\omega)$, where $\pi(s, g_\omega)$ is the universal policy.

State abstraction. A state abstraction is a mapping, say ϕ , that maps the original state space S to some finite lower-dimensional abstract space. If $\phi(s) = \phi(s')$ for a pair $s, s' \in S$, then the two states in the original state space are represented by the same state in the abstract space.

Reward abstraction. A reward abstraction means that the rewards given as part of the main task are hidden from the subtask policies and only received by the highest-level policy.

'Global' initiation condition. In the case of *definition of a subtask* in the tables provided in section 3, wherever the term 'global' is used in reference to the initiation condition I_ω , it means that the subtask can be chosen for execution in any of the states in the global state space.

3 APPROACHES FOR HIERARCHICAL REINFORCEMENT LEARNING

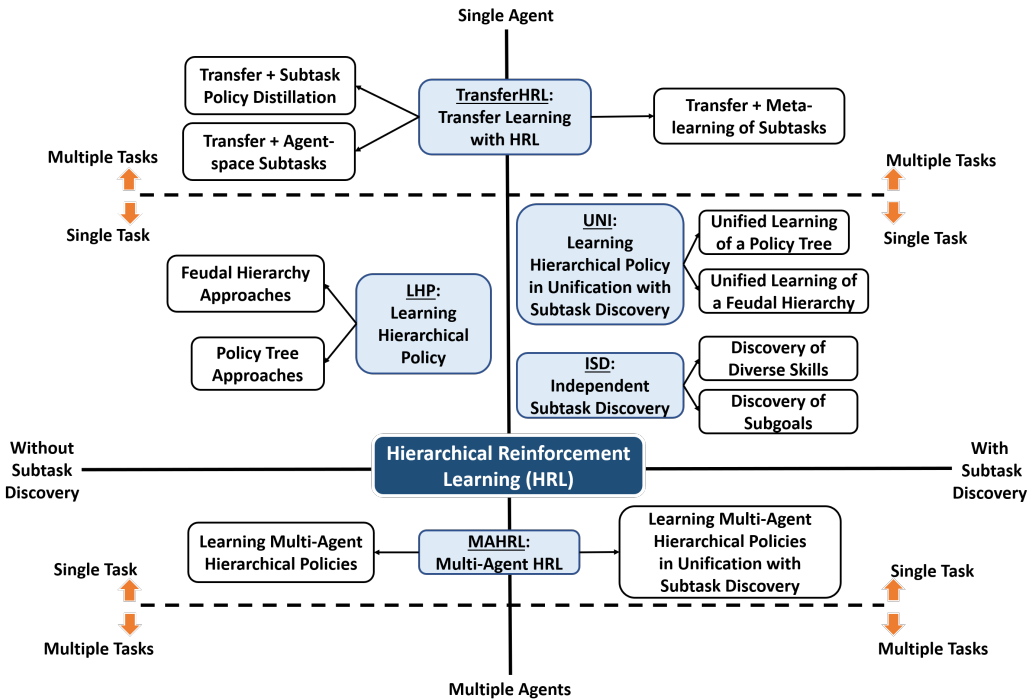


Fig. 3. **A taxonomy of HRL approaches.** The approaches are arranged along the following three dimensions: with or without subtask discovery, for single agent or multiple agents, and for single task or multiple tasks.

This section presents the review and classification of a wide range of approaches developed for Hierarchical Reinforcement Learning. The review is structured according to a novel taxonomy

depicted in Fig. 3. To design this taxonomy, the surveyed HRL approaches are arranged along three independent dimensions, which are as follows: 1. Approaches with subtask discovery or without it. 2. Approaches for training single agent or multiple agents. 3. Approaches for learning on single task or multiple tasks. There can be eight possible divisions (octants) of this three-dimensional space. We identify that the majority of the HRL approaches lie in six of these divisions. The approaches are further grouped into five major classes - LHP, UNI, ISD, MAHRL, and TransferHRL, based on the broad challenge addressed by them. The divisions and classes are as follows:

- (1) *Single agent, single task, without subtask discovery*: The approaches in this division are grouped into a major class called *Learning Hierarchical Policy* (LHP). The LHP approaches address the challenge of learning the hierarchical policy $\pi_{hierarchy}$ of an HRL agent without concerning with subtask discovery. They use handcrafted subtasks.
- (2) *Single agent, single task, with subtask discovery*: The approaches in this division are grouped into two major classes. The first class is called *Learning Hierarchical Policy in Unification with Subtask Discovery* (UNI). The UNI approaches address the challenge of learning the subtask space $\Omega_{hierarchy}$ and the hierarchical policy $\pi_{hierarchy}$ in a unified or end-to-end manner. The second class is called *Independent Subtask Discovery* (ISD). The approaches in this class address the challenge of discovering task-agnostic subtasks independently from any specific task. The subtask discovery is usually performed in a pre-training stage and then the subtasks are used to learn HRL agents on downstream tasks.
- (3) *Multiple agent, single task, without subtask discovery*.
- (4) *Multiple agent, single task, with subtask discovery*.

The approaches in divisions 3 and 4 are grouped into one class called *Multi-Agent HRL* (MAHRL). The MAHRL approaches broadly address the challenge of learning to coordinate among multiple HRL agents on a single joint task.

- (5) *Single agent, multiple tasks, without subtask discovery*.
- (6) *Single agent, multiple tasks, with subtask discovery*.

The approaches in divisions 5 and 6 are grouped into one class called *Transfer learning with HRL* (TransferHRL). The TransferHRL approaches broadly address the challenge of learning to transfer the hierarchical policy, subtasks, or other knowledge of an HRL agent across multiple tasks, where the subtasks may be handcrafted or discovered from scratch on multiple tasks.

We could not find considerable work in the remaining two divisions which represent the approaches for *multiple agents, multiple tasks, with/without subtask discovery*. A detailed discussion on the above-mentioned major classes, their associated challenges, and the approaches which address those challenges is provided in the following subsections. Before proceeding, we encourage the reader to recall the important definitions of the common terms provided in subsection 2.3.

3.1 Learning Hierarchical Policy (LHP)

This subsection reviews the approaches which address the challenge of learning the hierarchical policy $\pi_{hierarchy}$, without subtask discovery, for single agent and single task setting. Learning the hierarchical policy is a non-trivial challenge even when the subtasks are handcrafted. This is because any approach to learn a hierarchical policy must tackle the following key issues: carefully designing the algorithm to learn the policies at various levels of the hierarchy (including reward propagation, value function decomposition, state/action space design etc.), dealing with non-stationarity due to simultaneously changing policies, ensuring the optimality of the hierarchical policy as a whole, interpretability of the hierarchical policy, among other issues. Please refer to Table 1 for a summary of the LHP approaches.

The LHP approaches can roughly be put into two sub-classes. The first sub-class is called *feudal hierarchy*, in which the action space of a higher-level policy consists of subgoals corresponding to various subtasks. A subgoal chosen by the higher-level policy is taken as input by a universal policy³ at the lower-level. The objective of this lower-level universal policy is to achieve the input subgoal. The universal policy at each level can be treated as a sub-agent (as part of the HRL agent) which can perform all the possible subtasks at that level. This leads to the *feudal* concept of a "manager" sub-agent (higher-level policy) directing a "worker" sub-agent (lower-level policy) [5, 50, 76]. The corresponding approaches are discussed in subsection 3.1.1.

The second sub-class is called *policy tree*. This name is introduced for the purpose of classification in this survey, and it may not be found in the HRL literature. In a policy tree, the action space of a higher-level policy consists of the *different* lower-level policies of subtasks. The subtasks are not represented by a single universal lower-level policy, rather each subtask has a separate policy. The higher-level policy and the various lower-level policies form a 'tree'. The higher-level policy selects a lower-level policy directly for execution without intermediate subgoals [11, 12, 10]. The corresponding approaches are discussed in subsection 3.1.2.

3.1.1 Feudal Hierarchy Approaches. Dayan and Hinton [5] introduced the foundational feudal hierarchy called Feudal Reinforcement Learning (Feudal RL). Feudal RL is briefly described as follows: a higher-level manager sets a subtask which is to be executed by a lower-level worker. This is a pairwise relation such that the worker becomes the manager of the level below it, if the hierarchy has more than two levels. The manager communicates the subtask to the worker via a subgoal, where a subgoal is simply a state in the original or abstract state space. The objective of the worker is to reach the given subgoal. The worker at each level is a universal policy. The task reward may only be observed by the manager at the highest-level while the workers at other levels learn using the rewards for reaching the subgoals. The authors evaluate Feudal RL on a maze navigation task in a grid-world environment using predefined states at different spatial distances as subgoals. Feudal RL converges to a shorter path to a main goal state in the maze faster than the standard Q-learning. Rest of the feudal hierarchy approaches discussed in this paper are based on the concept of Feudal RL described above.

Kulkarni et al [50] proposed a *Deep HRL* approach consisting of a two level hierarchy of Deep Q Networks (DQNs) [41] which represent the manager and the worker. The manager network selects a subgoal from a set of predefined subgoals. It is learned using the task reward. The worker network takes the subgoal as input and selects primitive actions to achieve the subgoal. It is learned using rewards for reaching the subgoal. A subgoal is either an original state or an abstract representation of a state, for e.g. an object extracted from a image where the image represents the original state. One of the tasks used for evaluation is Atari Montezuma's Revenge in which standard DQN makes almost no progress in terms of collected rewards over 2 million time steps. On the other hand, Deep HRL with predefined subgoals progressively gets higher rewards as training proceeds.

Learning multiple levels of policies simultaneously results in **the issue of non-stationarity** [76, 91]. This means that the *state-action-next state* transition data, which is generated by executing a lower-level policy and is observed by a higher-level policy, varies over different time instances even for the same subgoal chosen in the same state of the environment. This is because the the lower-level policy is not stationary and its response to a (*state, subgoal*) pair changes during the learning process. Non-stationarity may result in many useless data samples and it needs to be addressed for data-efficient learning.

To address the issue of non-stationarity, Nachum et al [76] proposed a two-level feudal hierarchy with a mechanism for *subgoal re-labelling*. This approach is called Hierarchical Reinforcement

³defined in subsection 2.3.

Table 1. Approaches for Learning Hierarchical Policy (LHP).

Approach	Definition of a subtask ω					State abstraction	Reward abstraction	Suitable task domains		Main utilities	Main limitations
	I_ω	π_ω	β_ω	g_ω	r_ω			state space	action space		
Feudal Hierarchy Approaches											
<i>Feudal RL</i> ; Dayan & Hinton [5]	global	learned using r_ω	fixed time steps or at g_ω	PRE	w.r.t. g_ω	Possible	Yes	D, S	D	learns hierarchical policy using subgoals	hierarchical policy is likely to be sub-optimal // does not address the issue of non-stationarity
<i>Deep HRL</i> ; Kulkarni et al [50]	global	learned using r_ω	fixed time steps or at g_ω	PRE	w.r.t. g_ω	Possible	Yes	D/C, S/L	D	learns hierarchical policy using subgoals in the tasks with large and high-dimensional state spaces	hierarchical policy is likely to be sub-optimal // does not address the issue of non-stationarity
<i>HIRO</i> ; Nachum et al [76] // <i>HAC</i> ; Levy et al [91]	global	learned using r_ω	fixed time steps or at g_ω	PRE, using selective features of the state space	w.r.t. g_ω	Possible	Yes	C, L/S	C	reduce non-stationarity while learning a feudal hierarchy	manually mapping the state space to the subgoal space may be challenging if the state space is high-dimensional
<i>HAL</i> Jiang et al [89]	global	learned using r_ω	fixed time steps or at g_ω	PRE, using natural language instructions	w.r.t. g_ω	Possible	Yes	D/C, S/L	C	interpretability of hierarchical policy using natural language	requires predefined mapping of language instructions to the task domain
Policy Tree Approaches											
<i>Options</i> ; Sutton et al [11]	PRE	PRE, further tuned using task rewards	PRE	n.r.	optional	No	No	D, S	D	learns optimal hierarchical policies	all the Options are blended into the same MDP, so Options are not easily transferable
<i>MAXQ</i> ; Dietterich [12]	global	learned using r_ω	PRE	n.r.	PRE	Possible	Yes	D, S	D	policies of the sub-tasks are transferable across different tasks/hierarchies	hierarchical policy is likely to be sub-optimal
<i>HAMS</i> ; Parr & Russell [10]	PRE	learned using task rewards	PRE	n.r.	optional	No	No	D, S	D	prior knowledge about higher-level transition dynamics can be added by using Finite State Machines (FSMs) [69]	complicated to implement due to the intricate programming of FSMs

w.r.t. = *with respect to*; n.r. = *not required*.

PRE = predefined, D = discrete, C = continuous, S = small, L = large.

All the task domains are assumed to have moderate to extreme reward sparsity.

Learning with Off-policy Correction (HIRO). Subgoal re-labelling in HIRO can be described as follows: a two-level HRL agent interacts with its environment and gathers experience data. This data consists of transition tuples $(s_t, g_t, \sum r_{t:t+c}, s_{t+c})$ for the higher-level policy where r is the task reward, and $(s_{t+i}, g_t, a_{t+i}, r_{t+i}^{g_t}, s_{t+i+1}), \forall 0 < i < c$, for the lower-level policy. Here, c is a fixed time horizon for achieving each subgoal and r^{g_t} is the reward for achieving g_t . This data is subsequently used for training the HRL agent's hierarchical policy. However, if the agent does not achieve g_t after c time steps, then the subgoal is re-labelled in the transition data with another subgoal g'_t drawn from a distribution of subgoals which maximize the probability of the observed transitions. Then, the higher-level policy treats g'_t as its output in the hindsight, which correlates better with the observed transitions. In this way, HIRO reduces the effective non-stationarity. A subgoal in HIRO is defined by choosing selective features from a state in the original state space. HIRO is evaluated on MuJoCo [37] continuous control tasks and shown to perform better than standard RL and few other HRL approaches [67, 57].

Concurrently with HIRO, Levy et al [91] proposed an approach called the Hierarchical Actor-Critic (HAC) which also addresses the non-stationarity issue by subgoal re-labelling. In this scheme, the output subgoal in the higher-level data and the input subgoal in the lower-level data are replaced with the actual state achieved by the agent in hindsight rather than a probability based sampling of a new subgoal as in HIRO. This simple scheme also makes it possible to extend the hierarchy to more than two levels. HAC is evaluated on MuJoCo [37] continuous control tasks and shown to perform better than standard RL and HIRO. Moreover, the authors find that a three-level hierarchy performs better than a two-level hierarchy on those tasks.

A feudal hierarchy can also be realized by using *natural language instructions* instead of subgoal states. Jiang et al [89] proposed an approach called **Hierarchical Abstraction with Language** or HAL. They assume that a mapping from a state $s \in S$ to a language instruction⁴ $l \in L$ is given by a human supervisor or generated using a predefined program, as a probability distribution over the instructions in the instruction space L . On this basis, a Boolean function $\Psi : S \times L \rightarrow (0, 1)$ is defined. Now, if a higher-level policy selects an instruction l , then all the states for which $\Psi(s, l) = 1$ are the subgoal states for that instruction. In terms of the notation used in this paper, a language instruction l is a subtask ω and all the states which satisfy $\Psi(s, l) = 1$ belong to the subgoal set of ω , i.e. g_ω . The language instruction is encoded using a recurrent neural network [38] whose output is given to the lower-level policy as input. The subtask reward r_ω used to train the lower-level policy is a binary reward for achieving the states which satisfy the given instruction. HAL also handles non-stationarity using a mechanism called Hindsight Instruction Re-labelling. HAL is evaluated on object arrangement tasks in MuJoCo continuous control domain [37] integrated with CLEVR language dataset [49] and found to outperform standard RL and few other HRL approaches [56, 76].

3.1.2 Policy Tree Approaches. Sutton et al [11] introduced the Options framework in which the default MDP⁵ (called *core* MDP) is extended by including a set of subtasks, called *Options*, into the action space of an agent. A primitive action itself is considered as a single-step Option. An Option ω is defined using a tuple: $(\langle I_\omega, \pi_\omega, \beta_\omega \rangle)$. $I_\omega \subset S$ is a set of states in which the Option ω can be initiated or is valid. $\beta_\omega : S \rightarrow [0, 1]$ gives the probability of termination of ω in a state $s \in S$. $\pi_\omega : S \times A \rightarrow [0, 1]$ is the policy of the Option. An Option *may or may not* have an associated reward function r_ω . It is usually assumed that the Options are predefined by a programmer. The policy of an Option (π_ω) is predefined as an initial policy which can be fine-tuned in the context of a particular task using a mechanism called intra-option learning. Intra-option learning uses the task rewards even if the Option-specific reward r_ω is given. Moreover, the Q-value function of any Option ω , i.e. $Q_\omega(s, a)$, represents the expected cumulative reward until the end of the main task rather than only until the horizon (or termination) of the Option itself. Hence, an Option is not a standalone subtask unit but it *blends* into the core MDP. This limits the transfer-ability of the learned Options to other tasks but theoretically guarantees the optimality of the learned hierarchical policy. Sutton et al [11] evaluated the Options framework on grid-world tasks and found that it converges to optimal performance faster than standard RL.

MAXQ value function decomposition, proposed by Dietterich [12], decomposes the core MDP into smaller sub-MDP components. Each sub-MDP is associated to a subtask whose policy can be learned separately from other subtasks. This is achieved by decomposing the main Q-value function into the separate Q-value functions of the subtasks. Due to such decomposition, the Q-value of any subtask ω , i.e. $Q_\omega(s, a)$, represents the expected cumulative reward only until the horizon (or termination) of that subtask. Thus, each subtask policy can be learned as a standalone unit. This

⁴We use different notations compared to the original HAL paper [89] to avoid conflicts with the subtask notations (subsection 2.2.1).

⁵Markov Decision Process defined in subsection 2.1.

is in contrast against Options [11] in which a subtask is blended into the core MDP. The benefit of value decomposition is that the policies of the subtasks learned using MAXQ on one task can be easily transferred to different tasks or different MAXQ agents. The disadvantage of treating subtasks as standalone units is that the optimality of the hierarchical policy is compromised. MAXQ hierarchies are only *recursively* optimal, which means that MAXQ can learn optimal policies of the subtasks but the overall hierarchical policy is generally not optimal. In MAXQ, the predefined components of a subtask ω are its termination condition β_ω (e.g. events, conditions, subgoal states etc.) and reward function r_ω . A subtask may choose primitive actions in addition to choosing other subtasks. MAXQ is evaluated in a grid-world Taxi domain where it outperforms standard Q-learning.

Hierarchy of Abstract Machines (HAMs), by Parr and Russell [10], uses stochastic Finite State Machines (FSMs) [69] to represent the subtasks. HAMs is defined by a collection of stochastic FSMs $\{\mathbb{H}\}$. Each FSM $\mathbb{H}_i \in \{\mathbb{H}\}$ is defined using various machine-states. A machine-state is different from the state of the environment. An FSM consists of four *types* of machine-states: *action*, *call*, *choice*, and *stop*. There may be multiple machine-states of the same type. For e.g., Parr and Russell [10] illustrate an FSM with two *call* machine-states - 'Follow Wall' and 'Back Off' - which can be chosen from the *choice* machine-state of that FSM. The *choice* machine-state non-deterministically selects the next machine-state of \mathbb{H}_i . The possible choices consist of the various *action*, *call*, or *stop* machine-states of \mathbb{H}_i . A Q-table captures the values of the *(state, choice)* pairs. The choices are made using a greedy policy. Thus, reinforcement learning (specifically, Q-learning) is used to learn the optimal choice to be made in the *choice* state, given the current state of the environment. If an *action* machine-state is chosen, it executes a primitive action in the environment. If a *call* machine-state is chosen, it suspends the execution of \mathbb{H}_i and initiates the execution of another FSM, say \mathbb{H}_j . This is equivalent to a transition from one subtask to another. If a *stop* machine-state is chosen, it terminates the execution of \mathbb{H}_i and returns the control to the FSM that called \mathbb{H}_i . HAMs is a useful and elaborate framework for programming the prior knowledge about the transition dynamics in a domain. However, due to its intricate structure, HAMs has not found widespread application in HRL.

3.2 Learning Hierarchical Policy in Unification with Subtask Discovery (UNI)

This subsection reviews the approaches which unify subtask discovery with the learning of the hierarchical policy, for single agent and single task setting. For an HRL agent to be deployable on a particular task without requiring predefined or handcrafted subtasks, it is essential that the subtask space ($\Omega_{\text{hierarchy}}$) is discovered *simultaneously* when the hierarchical policy of the agent is being learned. This requires a unified or end-to-end learning approach which uses the same data for both subtask discovery and learning hierarchical policy, where the data is collected while performing the given task. A few important issues involved with such a unification are as follows: discovering the subtask space which maintains the optimality of the hierarchical policy, learning various components of a subtask during discovery (terminations conditions, initiation conditions, subgoals etc.) from scratch, discovering a dynamic number of subtasks, among other issues.

The UNI approaches discussed henceforth are based on the two types of LHP approaches - the feudal hierarchy (subsection 3.1.1) and Options (a policy tree approach, subsection 3.1.2). Please refer to Table 2 for a summary of the approaches.

3.2.1 Unified Learning of Policy Tree. The policy tree approaches discussed in this subsection discover Options [11] in unification with the learning of a hierarchical policy. As discussed in subsection 3.1.2, an Option is a subtask which is represented using three key components: initiation condition I_ω , policy π_ω , and a termination probability function β_ω . Unified Option discovery

Table 2. **UNI**: Approaches for Unifying Subtask Discovery with the Learning of Hierarchical Policy.

Approach	How are subtasks discovered?	Definition of a subtask ω					No. of sub-tasks	Suitable task domains		Main utilities	Main limitations
		I_ω	π_ω	β_ω	g_ω	r_ω		state space	action space		
Unified Learning of Policy Tree											
Daniel et al [45]	by learning a graphical model to segment trajectories into Options	global	learned using the task rewards	learned	n.r.	n.r.	PRE	D/C, S	D/C	learns trajectory-based Options from scratch	Options are not available during initial learning stage // number of Options needs to be predefined
<i>Skill Chaining</i> ; Konidaris et al [30]	by learning a classifier to identify initiation states of one Option ω and setting them as subgoals of another Option ω' , i.e. $g_{\omega'} = I_\omega$	learned using the classifier	learned using r_ω	at g_ω	learned using the classifier	w.r.t. g_ω	DISC	D/C, S	D	learns a flexible number of subgoal-seeking Options from scratch	Options are not available during initial learning stage // works only for goal-directed tasks
<i>Option-Critic</i> ; Bacon et al [56]	learned using policy-gradients derived using the task rewards	global	learned using the task rewards	learned	n.r.	n.r.	PRE	D/C, L/S	D	learns general Options from scratch, starting from initial learning stage	number of Options needs to be predefined // performs poorly on sparse reward tasks [67]
<i>Proximal Policy Option-Critic</i> ; Kliszarov et al [61]	learned using policy-gradients derived using the task rewards	global	learned using the task rewards	learned	n.r.	n.r.	PRE	D/C, L/S	C	extends Option-Critic to continuous action tasks	number of Options needs to be predefined // performs poorly on sparse reward tasks [76, 97]
<i>Hierarchical Option-Critic</i> ; Riemer et al [80]	learned using policy-gradients derived using the task rewards	global	learned using the task rewards	learned	n.r.	n.r.	PRE	D/C, L/S	D	extends Option-Critic to arbitrary number of levels of the hierarchical policy	number of Options needs to be predefined // might perform poorly on sparse reward tasks
<i>Interest-Option-Critic</i> ; Khetarpal and Precup [90]	learned using policy-gradients derived using the task rewards	learned	learned using the task rewards	learned	n.r.	n.r.	PRE	D/C, L/S	D	adds Interest Function to Option-Critic for learning the initiation condition I_ω	number of Options needs to be predefined // might perform poorly on sparse reward tasks
<i>Double Actor-Critic</i> ; Zhang et al [96]	learned using policy-gradients derived using the task rewards	global	learned using the task rewards	learned	n.r.	n.r.	PRE	D/C, L/S	D/C	learns general Options using off-the-shelf policy optimization algorithms instead of algorithms customized for Option-based SMDP	number of Options needs to be predefined // might degenerate to frequently terminating Options due to the lack of deliberation cost
Unified Learning of Feudal Hierarchy											
<i>FuN</i> ; Vezhnevets et al [67]	by learning the subgoal space using policy gradients	global	learned using r_ω	after fixed no. of steps or at g_ω	learned	w.r.t. g_ω	DISC	D/C, L/S	D	learning low-dimensional subgoal space	learned representation of the subgoal space may not be optimal
Nachum et al [77]	by learning the subgoal space using a sub-optimality bound	global	learned using r_ω	after fixed no. of steps or at g_ω	learned	w.r.t. g_ω	DISC	C, L/S	C	learning low-dimensional subgoal space, with theoretical guarantee of reduced sub-optimality	not expressive enough to represent subgoal-free behaviour (unlike Options), e.g. drive through traffic, move in a circle, running etc.

w.r.t. = *with respect to*; n.r. = *not required*.

D = discrete, C = continuous, S = small, L = large. PRE = predefined, DISC = discovered.

All the task domains are assumed to have moderate to extreme reward sparsity.

mainly involves the learning of π_ω and β_ω associated with each option ω simultaneously with the higher-level task policy (π_T). The initiation I_ω may also be learned or set as global⁶. None of these components are handcrafted.

Daniel et al [45] propose a graphical model for unified Option discovery. At any time-step t , the observed variables of this graphical model are the environment state s_t and the primitive action a_t . The hidden variables are the Option ω_t active at that time-step and the termination label b_t (binary, 0 or 1) indicating whether the previous Option ω_{t-1} terminates ($b_t=1$) or not ($b_t=0$). The model consists of the activation policy $\mu(\omega|s; \theta_O)$ which is parameterized using θ_O and provides the probability distribution over all the Options, conditioned on the current state s_t . Each Option ω consists of a sub-policy $\pi(a|s; \theta_A^\omega)$ and a termination function $\beta(b|s; \theta_B^\omega)$. The set of θ_A^ω parameters of all the Options is denoted as θ_A . Similarly, The set of θ_B^ω parameters of all the Options is denoted as θ_B . Then, the parameters $\{\theta_O, \theta_A, \theta_B\}$ are learned using Expectation-Maximization (EM) [3] with weighted estimates in the maximization step. The EM algorithm takes a set of trajectories as input and learns the parameters which best explain the trajectories. The trajectories are obtained from the exploration done by the HRL agent. In the maximization step, each state-action (s, a) pair is weighted in proportion to the advantage value $Q(s, a) - V(s)$. In this way, the Option learning process incorporates the expected task reward. Overall, the parameters $\{\theta_O, \theta_A, \theta_B\}$ for the activation policy, all the sub-policies, and all the termination functions are learned in conjunction during EM. Hence, this is a unified HRL approach. This approach is evaluated on grid-world navigation task with discrete state-action space and pendulum swing task with continuous state-action space. It is found to perform better than standard RL and few other subtask discovery approaches [17, 24]. The main limitation of this approach is that the number of Options need to be predefined and fixed.

Konidaris et al [30] propose an approach called *skill chaining* to incrementally construct Options while learning an HRL agent. The process of skill chaining starts with one Option ω created with its subgoal set (g_ω) containing the goal states of the main task. The Option policy π_ω is learned to reach the subgoals. A classifier is then learned to find the initiation states of ω . The states from which the subgoals of ω can be reached within a predefined number of time steps are classified as positive. In contrast, those states from which the subgoals cannot be reached within the time limit are classified as negative. The positive states are added to the initiation set I_ω . Now, another Option ω' is created, whose *subgoal set $g_{\omega'}$ is equal to I_ω* . Then, the above-mentioned classification and policy learning processes repeat for ω' , and so on, creating a chain of Options (or skills). Skill chaining discovers a flexible number of Options. The authors evaluate it on a Pinball control domain in which it outperforms standard RL. Recently, Bagaria et al [97] extended skill chaining for continuous control using deep reinforcement learning. Key limitations of skill chaining are that it works only on the tasks which have explicit goal states and requires strong initial exploration without Options to reach the goal states.

In all the approaches discussed above, an agent is not able to use the subtasks (Options) during the initial episodes of learning. This is because the agent must first collect experience trajectories to use for discovering the initial set of Options. Bacon et al [56] propose an Option framework which can learn the Options along with the entire hierarchical policy from the beginning of the learning process. This approach is called Option Critic (OC). In OC, a fixed number of Options are randomly initialized with parameterized policies and termination functions. The higher-level policy is also randomly initialized. Then, the higher-level policy and all the Options (policies and termination functions) are learned using policy gradients derived using the main task rewards. OC does not use subgoals or Option-specific rewards and it is theoretically guaranteed to learn optimal hierarchical

⁶Defined in subsection 2.3.

policy using the policy gradients. In evaluation, OC outperforms standard RL in Atari games like Seaquest, Ms Pacman etc using deep learning. It also performs better than skill chaining [30] in the initial stages of learning in the Pinball domain. A limitation of OC is that it requires the *number* of Options to be predefined. Moreover, the policy gradients are heavily dependent on the main task reward. Thus, OC is likely to perform poorly on the sparse-reward tasks [67, 76]. OC has become a popular framework for unified HRL due to its strong theoretical foundation. This has led to the emergence of various approaches that build upon OC [61, 80, 90], which are listed in Table 2.

The Option-Critic (OC) framework requires the policy optimization algorithms which are customized for the Option-based Semi-Markov Decision Process (SMDP) [11]. This puts a restriction on using other advanced policy optimization algorithms [43, 72] for HRL in an off-the-shelf manner. To address this issue, Zhang et al. [96] proposed the formulation of an HRL hierarchy as two parallel *augmented* Markov Decision Processes (MDPs) [35], where each MDP uses an augmented state space that is a cross-product of the original state space and the set of Options. The higher-level MDP models the problem of learning a policy over Options and their corresponding termination conditions, while the lower-level MDP models the problem of learning the Option policies. Both MDPs use the task rewards only, without subgoals or subtask rewards. The proposed approach is called Double Actor Critic (DAC) because it applies actor-critic algorithms to learn the policies at both the levels. Due to the augmented MDP formulation, DAC is compatible with various off-the-shelf actor-critic algorithms for policy optimization. It shows significantly better performance than OC on various robot simulation tasks. However, similar to OC, DAC requires the number of Options to be pre-defined.

It has been observed that OC and DAC approaches might encounter a higher-level policy *degeneration* phenomenon [74, 96] when the probability of termination of each Option becomes very high after a long period of training. This causes the higher-level policy to switch from one Option to another at almost every time step and the learned Options do not specialize on any recognizable behaviour. Harb et al. [74] argue that temporally extended Options are not necessary from the theoretical perspective of policy optimization, because the optimal policy is actually achieved in terms of the primitive actions. To learn temporally extended Options, they introduced a regularizer called *deliberation cost*, which is a penalty received by the higher-level policy upon switching Options. This encourages the higher-level policy to retain each Option for a longer period of time, which leads to an *empirically* better performance than the frequently switching Options. The deliberation cost is included in the OC framework but the authors of DAC leave the integration of such a cost for future work. Another form of degeneration might occur when the higher-level policy selects a single Option for the entire task duration. This is possibly because each Option can be initiated anywhere in the state space if it uses a global initiation condition I_ω . Khetarpal and Precup [90] proposed an approach called Interest-Option-Critic (IOC), built on top of the OC framework, which learns the initiation condition of each Option using the policy gradients so that various Options can specialize on different behaviours, thereby preventing the degeneration.

3.2.2 Unified Learning of Feudal Hierarchy. In a feudal UNI hierarchy, the subgoal space and the universal policies at all the levels are learned in a unified manner.

Vezhnevets et al [67] propose a feudal hierarchy of neural networks in which a higher level network called the "Manager" samples a subgoal in a *learned* latent subgoal space. The subgoal may be a point in the latent space or it may be a unit vector representing a direction in the latent space. The subgoal is taken as input by a lower-level network called the "Worker" which must learn a policy to achieve the subgoal using the distance to the subgoal as a reward. The Worker is trained using the usual policy gradients derived from the subgoal-based rewards. The Manager is trained using a *transition gradient* introduced by the authors. This gradient is derived using the

task reward as well as the distance between the subgoal assigned to the Worker and the actual state-transition made by the Worker. Thus, the Manager learns from both the task rewards and the Worker’s behaviour. The transition gradient is used to learn both the Manager policy and the latent subgoal space. This approach, called Feudal Networks (FuN), shows better performance on Atari games, such as Montezuma’s Revenge, than DQN [41] and Option-Critic [56].

Feudal Networks do not guarantee that the learned subgoal space leads to optimal hierarchical policy. Nachum et al [77] address this issue in a subgoal representation learning approach developed upon HIRO [76] (HIRO is discussed in subsection 3.1.1). The authors derive an optimization objective based on a theoretical bound on the sub-optimality of the hierarchical policy. This objective is used to learn a function $f_\theta(s)$ which transforms the state space to a low-dimensional subgoal space. Thus, the learned subgoal space representation minimizes the sub-optimality and the hierarchical policy based on HIRO addresses the non-stationarity issue. This approach is evaluated on MuJoCo continuous control tasks [37], using both low-dimensional state spaces and high-dimensional state spaces (e.g. images used as states). It outperforms various other subgoal representation schemes such as direct use of the original states (e.g. images) as subgoals, using latent space learned in the style of FuN [67], using subgoal embedding derived from Variational Autoencoder [39] etc.

3.3 Independent Subtask Discovery (ISD)

This subsection reviews the approaches for independent subtask discovery, for single agent and single task setting. The ISD challenge arises when the aim is to automatically find the subtasks which are task-agnostic and transferable to HRL agents across various unknown tasks. The implication is that the process of subtask discovery should be *independent* of the process of learning the hierarchical policy of an agent. The common idea behind the ISD approaches is similar to that of pre-training in the general context of machine learning [33]. In an ISD approach, subtask discovery typically occurs in a pre-training stage. The discovered subtasks are then used to learn a hierarchical policy to perform a particular target task. This pre-training approach leads to a few key issues which are part of ISD, as follows: ensuring that the subtask discovery process is data-efficient, discovering subtasks which allow diverse exploration of the state space independent of any specific task, learning continuous subtask space which allows generalization through sampling of new subtasks on target tasks, among other issues.

It is important to make a *disclaimer* that we do not strongly classify a particular approach as ISD without any possibility of being applicable for subtask discovery in a unified manner. A subtask discovery algorithm provided as part of an ISD approach may also be used for unified learning by running it simultaneously with the training of the agent’s hierarchical policy, using the same task as the source of data for both. This is mentioned wherever applicable in the following discussion. Please refer to Table 3 for a summary of the ISD approaches.

3.3.1 Discovery of Subgoals. A set of subtasks can be constructed by discovering their corresponding subgoals (one of the objective components mentioned in subsection 2.2.1) and then learning one subtask policy to reach one subgoal or a set of related subgoals.

A straightforward approach for subgoal discovery is to find the bottlenecks in the state space of a task domain. **Bottlenecks** are those states at which several paths connecting various initial states to various goal states converge. Removing the bottlenecks is likely to cut access to the possible goal states in the state space. Hence, the discovered bottlenecks are treated as important subgoals and policies of subtasks are learned to reach them.

McGovern and Barto [15] proposed a frequency-based approach for bottleneck discovery. This approach uses the idea of Diverse Density [9] to find interesting bottlenecks. An agent explores the state space using an initial policy and collects the trajectories from various initial states to

Table 3. **ISD: Approaches for Independent Subtask Discovery.**

Approach	How is a subtask discovered?	Suitable task domains		Main utilities	Main limitations
		state space	action space		
Discovery of Subgoals					
McGovern et al [15] // Şimşek et al [29]	using frequently occurring bottlenecks* as subgoals	D, S	D	simple to implement frequency-based heuristic	might require excessive exploration to collect multiple trajectories for estimating the bottleneck frequencies
<i>Q-cut</i> ; Menache et al [17]	using graph-cut bottlenecks as subgoals	D, S	D	data-efficient bottleneck discovery using graphs over the global state-action space	not easily scalable to continuous or large state spaces
<i>L-cut</i> ; Şimşek et al [24]	using graph-cut bottlenecks as subgoals	D, L/S	D	data-efficient bottleneck discovery using local sub-graphs	not easily scalable to continuous state spaces
Macahdo et al [63]	using subgoals at local maxima of Proto Value Functions (PVFs) [27]	D/C, L/S	D	learning general subtasks for diverse exploration	computing PVFs may be computationally expensive
<i>HASSLE</i> ; Bakker et al [22]	using state space cluster centroids as subgoals	D, S	D	learning general subtasks for diverse exploration	not easily scalable to high-dimensional state spaces
<i>HSP</i> ; Sukhbaatar et al [82]	using a continuous latent embedding of subgoals learned via asymmetric self-play [65]	D/C, S/L	C	discovering continuous subgoal embedding instead of discrete subgoal space	cannot generalize to subgoal-free skills e.g. drive through traffic, move in a circle, running etc.
Discovery of Diverse Skills**					
<i>VALOR</i> ; Achiam et al [68]	by learning the policy π_ω to maximize the MI*** between the skill ω and the trajectories generated by executing it	D/C, S/L	D/C	discovers diverse skills and provides a procedure to grow the number of skills	requires excessive exploration data to maximize MI
Hausman et al [75]	by learning the policy π_ω to minimize the cross-entropy between the skill embedding and the generated trajectories	D/C, S/L	C	discovers continuous embedding of skills instead of a discrete set	requires excessive exploration to learn both the skill embedding and the policy
<i>SecTAR</i> ; Co-Reyes et al [79]	by learning the policy π_ω to decode the latent embedding of the trajectories correlated with the skill ω	D/C, S/L	D	discovers continuous embedding of skills and additionally provides a technique for unified learning	requires excessive exploration to collect diverse trajectories for embedding

*defined in subsection 3.3.1; **definition of 'skill' is given in subsection 2.3; ***Mutual Information w.r.t. = *with respect to*; n.r. = *not required*.

D = discrete, C = continuous, S = small, L = large. PRE = predefined.

various terminal states. Each trajectory is classified as *positive* or *negative* depending on whether its terminal state is a desired goal state or not. Here, a desired goal state is related to a pre-training task. Then, a Diverse Density (DD) measure is used to estimate the probability of a state occurring on the positive trajectories vs the probability of occurrence on the negative trajectories. The states with higher DD are chosen as the bottleneck subgoals. Simsek and Barto [29] also proposed a frequency-based bottleneck discovery approach which may be more data efficient because they build partial state-action transition graph from both the successful and unsuccessful experiences. In this graph, each node is a state and each edge is an action. The graph can be used to *simulate* new trajectories by simply following the connections between various nodes. The authors use a measure called *betweenness* defined as $\sum_{s \neq t \neq v} \frac{\sigma_{st}(v)}{\sigma_{st}} w_{st}$, where σ_{st} is the total number of optimal paths from a node s to a node t and $\sigma_{st}(v)$ is the number of such paths passing through a node v .

w_{st} is the weight of a (s, t) pair. If the betweenness score for v is higher than its neighbours by a certain threshold, then v is chosen as a subgoal.

Frequency statistics require a large number of trajectories to be generated in order to estimate the occurrence frequencies of various states. A more efficient approach is to simply construct a state-action transition graph once, in which each node is a state and each edge is an action, and then find the bottlenecks by analyzing the connectivity between different sub-graphs. Menache et al [17] proposed an approach called *Q-cut* which builds a state-action transition graph over the global environment and then finds bottleneck nodes using Max-Flow/Min-Cut algorithm [21]. Each bottleneck is treated as a subgoal for a subtask, and the subtask's policy is learned to reach it. Q-cut is not be easily scalable to large and/or continuous state spaces due to the need for constructing the *global* graph which requires diverse exploration. Şimşek and Barto [24] proposed a similar graph-cut approach to find the bottlenecks, but they use local sub-graphs rather than a global graph constructed over the entire state space. This is called *L-cut*. L-cut can find local bottlenecks incrementally and can be applied on large state spaces because an agent does not need to build the global graph. However, it can not be easily scaled to continuous state spaces because the graphs used by the algorithm require discrete states as nodes.

Bottleneck discovery may not be suitable for task domains which do not necessarily contain identifiable bottlenecks. This may happen when the paths to various desired goal states are diversely distributed and do not sufficiently converge on particular bottlenecks. Machado et al [63] proposed an approach to find subgoals which are distributed over diverse regions of the state space rather than only at the bottlenecks. This approach uses Proto Value Functions (PVFs)⁷ [27]. The subgoals lie at the local maxima of different PVFs defined over the state space. The PVFs are derived from the eigendecomposition of the Laplacian matrix⁸ of the state transition graph. Bakker & Schmidhuber [22] proposed unsupervised clustering directly in the state space in order to discover subgoals over diverse regions. An agent explores the state space and forms clusters of states using an online Adaptive Resource Allocation Vector Quantization (ARAVQ) algorithm [13]. The centers of various clusters are used as subgoals. The exploration and clustering process is not dependent on any particular task, hence the subgoal discovery is task-agnostic or independent. The authors also provide an algorithm to learn a hierarchical policy using the already discovered subgoals. In this algorithm, subgoals constitute the action space of a higher-level policy which is learned using the task rewards. The lower-level consists of multiple policies where each policy can specialize to reach of a subset of the subgoal space and learn using the rewards for subgoal achievement. A mapping function is learned to assign a subgoal to the appropriate lower-level policy. There is also a provision for subgoal re-labelling, similar to HIRO [76] or HAC [91] (subsection 3.1.1), for handling non-stationarity. The complete approach is called Hierarchical Assignment of Subgoals to Sub-policies LEarning (HASSLE).

A significant limitation of all the subgoal discovery approaches reviewed so far is that the discovered subgoal space is *discrete*. In such a discrete space, it is not feasible to interpolate new subgoals, thereby limiting the range (or diversity) of the tasks which can be performed using the discovered subgoal space. **Discovery of continuous subgoal spaces** is important to address this issue. Sukhbaatar et al [82] proposed an approach called Hierarchical Self Play (HSP) to learn continuous embedding of subgoals using asymmetric self-play [65]. Asymmetric self-play is an unsupervised pre-training phase which is illustrated as follows: at the beginning of asymmetric self-play, two standard (non-hierarchical) RL policies are initialized, say Alice (π_A) and Bob (π_B). Let the initial state be s_0 . First, Alice takes T_A steps in the environment starting from s_0 . Let $s^* = s_{T_A}^A$

⁷PVFs are the basis functions of the state transition graph of an MDP.

⁸<https://mathworld.wolfram.com/LaplacianMatrix.html>

be the final state reached using Alice. Next, the environment is reset to s_0 and s^* is given as a target to Bob. At each time step, s^* is encoded into a low-dimensional latent subgoal embedding using a learned encoder E , i.e. $g_t = E(s_t^B, s^*)$ where s_t^B is the state at time t while Bob is executing. The Bob policy selects primitive actions as $a_t^B = \pi_B(s_t^B, g_t)$. Bob is given T_B steps to achieve the target state s^* , otherwise it fails. It is given a reward $R^B = 1$ if the target is achieved, otherwise 0. Alice is given a reward $R_A = 1 - R_B$. The authors choose $T_B \sim T_A$. This results in an asymmetric self-play mechanism in which Bob learns to reach targets set by Alice and Alice learns to achieve new, unexplored targets which are beyond the reach of Bob. In this way, an agent consisting of these two policies explores the environment effectively without external supervision.

The subgoal encoder of HSP, $E(s_t^B, s^*)$, basically encodes various target states (s^*) into a low-dimensional subgoal space. Once a subgoal space has been discovered in the pre-training phase of asymmetric self-play, it is simply used as the *continuous* action space of a higher-level policy of an HRL agent used to perform a particular task. The lower-level of the agent is initialized using the Bob policy and then fine-tuned on the task.

3.3.2 Discovery of Diverse Skills. The subgoal discovery approaches are unsuited to find subtasks without definite subgoals associated with them. Consider an example of a subtask 'drive through traffic' which is a part of a longer horizon task of reaching a destination. This subtask requires an agent to manoeuvre a vehicle around traffic *without* any particular subgoal. Subgoal discovery is inapt for finding such a subtask. Therefore, general approaches are desired which can directly discover a set of diverse skills, instead of learning them through subgoals. As defined in subsection 2.3, a 'skill' refers to the policy of a subtask in the sense that it semantically represents the *ability to do something well*. After discovery, the skills can be added to an HRL agent as the lower-level policies and fine-tuned on a particular task.

One strategy for discovering diverse skills is by maximizing the Mutual Information (MI)⁹ between a skill ω and the states reached or the trajectories generated by using that skill. This is achieved by using a latent skill vector, say z_ω , and defining the skill policy as $\pi_\omega(s) = \pi(s, z_\omega)$, where $\pi(s, z_\omega)$ is a universal policy which takes the skill vector as input. z_ω is typically a one-hot encoded vector in most of such approaches [57, 47, 68, 71]. Maximizing the MI between z_ω and the states reached or trajectories generated by following $\pi(s, z_\omega)$ results in the discovery of a set of diverse skills correlated with diverse terminal states or trajectories. An overview of the specific approaches based on this strategy is provided below.

Florensa et al [57] proposed an approach called SSN4HRL in which the universal policy $\pi(s, z_\omega)$ is represented using a stochastic neural network [23]. This policy is learned by maximizing an MI reward $P(z_\omega | (x, y))$, where (x, y) is the location of the agent in a spatial environment and $P(z_\omega | (x, y))$ is the probability of predicting z_ω by observing (x, y) . This probability can only be maximized if the locations reached by executing different skills are sufficiently different. Gregor et al [47] proposed Variational Intrinsic Control (VIC) which discovers the skills by maximizing an objective containing an MI term $P(z_\omega | s_0, s_T)$ between the skill vector z_ω and the terminal state s_T reached by the policy $\pi(s, z_\omega)$, conditioned on the start state s_0 . Eysenbach et al [71] proposed an approach called Diversity is All You Need (DIAYN) which discovers the skills by optimizing an objective function designed to maximize the MI between a skill vector z_ω and every state in a trajectory generated by $\pi(s, z_\omega)$, without considering the order of the states in the trajectory.

All these approaches can discover a set of diverse skills but also share few common limitations, such as, they can only discover a predefined number of skills and their MI objectives either ignore trajectories or do not consider the order of states in the trajectories. To address these limitations,

⁹Mutual information (MI) of two random variables is the amount of information (in units such as *bits*) obtained about one variable by observing the other variable (<https://people.cs.umass.edu/~elm/Teaching/Docs/mutInf.pdf>).

Achiam et al [68] proposed Variational Autoencoding Learning of Options by Reinforcement (VALOR). In this approach, a policy $\pi(s, z_\omega)$ is basically an *encoder* which encodes z_ω into a trajectory. The trajectory is denoted as τ . It is fed into a *decoder* which must map it back to z_ω . This leads to an optimization objective aimed at maximizing $P(z_\omega|\tau)$. Thus, VALOR discovers a set of diverse skills which essentially correlate to a diversity of trajectories taken by an agent. VALOR also provides a procedure to discover a progressively increasing number of skills.

The ISD approaches discussed so far can only discover a *discrete* set of skills. A possible reason may be that computing the probability distributions over a continuous space of skills, for estimating the MI, is complicated. However, discovering a continuous space of skills is important for generalization, such that new skills can be easily interpolated in such a continuous space. Co-Reyes et al [79] proposed an approach called Self Consistent Trajectory Autoencoder (SeCTAR), in which an encoder LSTM [6] embeds the *state transition trajectories* (i.e. $\{s_t, s_{t+1}, s_{t+2} \dots\}$) into a low-dimensional continuous latent vector space and a decoder LSTM learns to decode a latent vector into a policy. A latent vector represents similar trajectories, hence, the policy decoded from the latent vector is considered to represent a skill. A diverse latent space is learned by using an exploration mechanism which generates diverse trajectories. This encoder-decoder model is learned in a pre-training phase for skill discovery. After the pre-training phase, the encoder module is removed and the encoded latent space is available as a continuous space of skills. A higher-level policy of an HRL agent samples a skill in this continuous space, which is then decoded by the decoder module in the form of a policy. In addition to independent skill discovery, the authors also extend SeCTAR for *unified learning* on a particular task. This is achieved via an *explorer policy* which explores the task relevant states and their neighbourhood such that the trajectories encoded by the SeCTAR encoder module are correlated to the given task. This way, the discovered space of skills is situated in the task on which the HRL agent is being trained concurrently.

Hausman et al [75] proposed a variational inference approach to learn diverse skills along with a continuous latent embedding of the skills. The objective used for learning is the minimization of the cross-entropy¹⁰ between a distribution over the latent embedding conditioned on a one-hot skill vector and a distribution over the latent embedding conditioned on the trajectory generated by executing that skill. The learned latent embedding is the continuous space of skills in which new skills can be interpolated.

3.4 Transfer Learning with HRL (TransferHRL)

This subsection reviews the approaches for learning single HRL agent on multiple tasks via multi-task or transfer learning. The TransferHRL class represents all such approaches with or without subtask discovery. The concept of transfer learning with HRL arises from the idea that certain subtasks are shareable across multiple related tasks. The policy learned to perform a subtask in the context of one task can be used to accelerate the adaptation of the agent on other related tasks as well. Therefore, this is an important challenge to address in HRL.

Transfer learning is used in the context of RL by transferring experience data, action policy, or Q-value function from one task to another task, in order to accelerate learning on the latter task [54, 36]. In the context of HRL, the problem of transferring the main task policy (the highest level in the hierarchical policy) can be addressed using the RL transfer approaches. However, the subtasks become the additional transferable components, including their policies, termination/initiation conditions, subgoals etc. This poses additional issues which are unique to the transfer learning with HRL (TransferHRL) compared to the transfer learning with standard RL. A few of the issues unique to TransferHRL are as follows:

¹⁰Cross-entropy is a measure of the difference between two probability distributions for a given random variable.

- Efficiently scaling to a large number of subtasks during lifelong transfer learning.
- Transferring subtasks across task domains with different state spaces.
- Learning subtasks from scratch on multiple related tasks.

These issues are discussed henceforth in this subsection, along with the respective approaches to address them. We reiterate that the ISD approaches (subsection 3.3) learn task-agnostic subtasks which can be transferred to various tasks. However, ISD approaches do not address the above-mentioned issues.

3.4.1 Transfer + Subtask Policy Distillation. An HRL agent might learn and accumulate a large number of subtask policies when transferring from one task to another, such as during lifelong learning [70]. Naively adding numerous subtasks to the action space of the higher-level policy may cause memory-inefficiency due to the need to store multiple subtask policies. Tessler et al [66] proposed a deep HRL framework, called **Hierarchical Deep Reinforcement Learning Network (H-DRLN)**, for memory-efficient transfer and retention of a set of pre-trained policies corresponding to various subtasks (without subtask discovery). These policies are pre-trained using manually defined objectives. The policies are transferred to an HRL agent (in a target task) via *multi-skill distillation* introduced by the authors. Multi-skill distillation is a form of policy distillation [54] which enables H-DRLN to efficiently combine the policies of several subtasks into a single distilled policy. This makes H-DRLN memory-efficient. When performing a target task, H-DRLN agent uses a higher-level policy to choose a subtask which is executed by the distilled policy (at the lower-level). H-DRLN learns to perform multiple tasks in the Minecraft game¹¹ with lower sample complexity and better performance compared to standard RL. The main limitation of H-DRLN is that the policies of subtasks need to be pre-trained using manually defined objectives (before distillation). Moreover, H-DRLN enables distillation of the existing policies into a target policy but does not provide a method for continual discovery of new subtasks.

3.4.2 Transfer + Agent-space Subtasks. A subtask policy learned in a particular task domain is conditioned on the state space S of that particular domain. If that policy is transferred to a new task domain whose state space is different from S , then the policy may no longer be optimal. To address this issue, it is necessary to achieve domain-invariant transfer such that the subtask policies are robust to the varying state spaces. Konidaris et al [26] proposed an approach to learn Options which can be transferred to various tasks with different state spaces. An Option is a subtask representation as defined in subsection 3.1.2. The invariance is achieved by using two separate state representations, one in the task-space and the other in the agent-space. The agent-space representation contains features which are dependent on the agent and invariant across various tasks. Option policies are learned in the agent-space and they can be transferred to various tasks with different task-spaces as long as the agent-space remains the same across those tasks. The higher-level policy, which selects an Option for execution, is learned using the task-space and adapts according to the active task. The authors design multiple versions of their own task domains to evaluate this method. It is found that introducing agent-space Options improves the agent's performance compared to regular Option-based techniques [11]. The main limitations of this approach are that the Options are not automatically discovered since their policies are learned using pre-defined rewards and the distinction of the agent-space from task-space needs to be manually defined.

3.4.3 Transfer + Meta-learning of Subtasks. Apart from efficiently transferring pretrained subtasks, it is also important to discover subtasks across multiple tasks from scratch. Frans et al [60] proposed a deep HRL approach for discovering Options via meta-learning on multiple related tasks.

¹¹<https://github.com/vkurenkov/hierarchical-skill-acquisition>

This approach, called Meta Learning Shared Hierarchies (MLSH), contains a set of Options with parameters shared across different tasks, which are learned from scratch by jointly maximizing the expected rewards obtained in all the tasks (meta-learning). MLSH does not use any other handcrafted rewards or subgoals. The learned Options are transferred to the unseen tasks by using task-specific higher-level policies. MLSH is evaluated in the MuJoCo [37] locomotion domain on which it outperforms standard RL and Option-Critic [56]. The main limitation of MLSH is that the number of Options needs to be fixed, which limits the diversity of tasks which can be performed using MLSH. Simply increasing the number of Options would require an agent to store and learn a large number of Option policies which is likely to be inefficient in terms of memory and computation.

3.5 Multi-agent Hierarchical Reinforcement Learning (MAHRL)

This subsection reviews the approaches for learning to coordinate multiple HRL agents on a single joint task. The MAHRL class represents all such approaches with or without subtask discovery. In the standard Multi Agent Reinforcement Learning (MARL) [7, 46, 58, 62], multiple agents learn to coordinate their primitive action policies by optimizing a joint task objective and, occasionally, by sharing common information about their states and actions. Going beyond standard MARL, the benefits of task decomposition can also be exploited for multi-agent coordination by dividing a joint task into multiple subtasks distributed across different HRL agents. Each agent can perform more than one subtask and different agents learn to coordinate their higher-level policies (that choose the subtasks) rather than just the primitive action policies. In certain cases, the agents may only coordinate at the higher-level while treating their primitive action policies as independent of each other [53]. Learning to coordinate among HRL agents is called Multi-agent HRL (MAHRL) [25].

Developing MAHRL agents is a non-trivial challenge due to different complex issues involved. All the usual issues found in the standard MARL also apply to MAHRL, such as the non-stationarity due to multiple agents learning simultaneously, partial observability due to limited communication, and difficulty of reward distribution in case of heterogeneous agents. However, the use of hierarchical policies (for all the agents) also introduces a few additional issues which are unique to MAHRL, described as follows:

- **Synchronization of subtask terminations across different agents.** The synchronization of the decisions made by different agents is affected by the terminations of their subtasks (i.e., β_ω) [18]. The subtasks of all the agents can be terminated in a *synchronized* manner by forcefully terminating all of them when the subtask of any one agent terminates or introducing a waiting time until all the subtasks terminate naturally. This scheme makes it easy to coordinate decisions but also puts a forced bias on the termination conditions which may cause sub-optimal behaviour. Another scheme is *asynchronous* termination in which each agent's subtask terminates independently of the terminations of other agents' subtasks. This scheme is free from any forced interruption but makes it difficult to coordinate decisions as they occur asynchronously.
- **Subtask discovery in the context of other agents' behaviours.** An agent must account for the presence of other agents in the environment during subtask discovery. Since other agents are non-stationary elements of an agent's environment, the discovered subtask space itself may become non-stationary. This also makes independent subtask discovery (ISD) quite challenging because the applicability of the discovered subtasks of an agent depends not only on the target task dynamics but also on the other agents present in a target task.

Definitions of key terms and concepts. A few important terms and concepts used in this subsection, *in the context of HRL*, are defined below.

Centralized agents implies a paradigm in which all the agents share the highest-level of their hierarchical policies using a single central policy (e.g. a central manager policy [86]).

Decentralized agents with centralized learning implies a paradigm in which all the agents have their own separate hierarchical policies (no central policy), but they learn by centralized information sharing (or communication). The information generally consists of the state observations and the subtasks being executed by different agents.

Fully decentralized agents implies a paradigm in which all the agents have decentralized hierarchical policies and learn *without* information sharing or central policies, under partial observability¹².

Homogeneous agents are those agents which have similar subtasks spaces and hierarchical structures. This implies that each agent may have the capability to independently perform the task and there may not be significant inter-dependencies among different agents.

Heterogeneous agents are those agents which have different subtask spaces and hierarchical structures. Hence, they are different in their capabilities and no individual agent may perform the entire task on its own. The agents may have strong inter-dependencies such that one agent assists another agent or it is impossible for an agent to receive task rewards unless another agent finishes its subtasks.

The key MAHRL approaches are discussed henceforth. Please refer to Table 4 for a summary of the approaches.

3.5.1 Learning Multi-Agent Hierarchical Policies. Ghavamzadeh et al [25] proposed an approach for cooperation among homogeneous agents at different levels of their task decomposition hierarchies. This approach is called *Cooperative HRL*. The authors use MAXQ [12] to define and learn the hierarchical policy of each agent. This approach is developed upon an earlier work by Makar et al [14]. The subtask space for each agent is predefined, using handcrafted reward function for each subtask. The task decomposition hierarchy of each agent may consist of more than two levels and each level can be defined as a cooperative or independent level. The agents cooperate in terms of their subtasks. The primitive action level is independent for all the agents. Moreover, the subtasks of different agents terminate using predefined *asynchronous* events. This approach uses *decentralized agents with centralized learning* paradigm. All the agents are homogeneous and they learn their decentralized hierarchical policies by sharing information about their local states and chosen subtasks. The authors report better performance of Cooperative HRL against standard MARL in an Automated Guided Vehicle (AGV) task for warehouse management and a Trash collection task. In order to ease the requirement of continuous information sharing, the authors also proposed a variant of this approach with communicate and non-communicate decisions at the higher levels of every agent's hierarchy.

Cooperative HRL addresses the issues of coordination and communication among homogeneous agents which generally do not have strong inter-dependencies. However, this approach is not trivial to apply in heterogeneous setting in which various agents perform different parts of a task and typically have strong inter-dependencies, such that few agents can reach the rewarding states while others take a preceding role to enable such agents. This implies that the task rewards may be too sparse or delayed for certain agents, making the learning process difficult. Engineering agent-specific rewards can be daunting in complex domains. Another limitation of the Cooperative HRL is that it uses termination events which need to be manually specified for all the subtasks.

¹²Entire information about the task environment, including the activity of other agents, is not available to an agent.

Table 4. MAHRL: Approaches for Multi Agent Hierarchical Reinforcement Learning.

Approach	Paradigm	How is a subtask discovered?	Definition of a subtask ω					Suitable task domains			Main utilities	Main limitations
			l_ω	π_ω	β_ω	g_ω	r_ω	state space	action space	type of agents		
Learning Multi-Agent Hierarchical Policies												
Cooperative HRL; Chavamzadeh et al [25]	decentralized agents with centralized learning	-	PRE	learned using r_ω	PRE, <i>async</i>	n.r.	PRE	D, S	D	HOM	MAHRL for homogeneous agents, with flexible communication	restricted to homogeneous agents
ISEMO; Pateria et al [94]	decentralized agents with centralized learning	-	PRE	PRE	learned, <i>async</i>	n.r.	n.r.	D/C, S/L	D	HOM/HETR	MAHRL for heterogeneous or homogeneous agents, using a reward augmentation strategy based on inter-dependencies among the agents	the initiation conditions of various subtasks must be predefined to capture the inter-dependencies among the agents
Tang et al [84]	decentralized agents with centralized learning	-	global	learned using r_ω	PRE, <i>sync or async</i>	n.r.	PRE	D/C, S/L	D	HOM	MAHRL for homogeneous agents using deep learning	restricted to homogeneous agents
FMH; Ahilan et al [86]	centralized agents	-	global	learned using r_ω	PRE, <i>sync</i>	PRE	w.r.t. g_ω	D/C, S/L	D	HOM/HETR	MAHRL using centralized control	learning the central high-level policy is likely to be complex if the joint subtask space is large
PoEM; Liu et al [53]	fully decentralized agents	-	PRE	learned from pre-existing demonstrations	PRE, <i>async</i>	n.r.	n.r.	D, S	D	HOM/HETR	MAHRL with full decentralization and under partially observability	requires preexisting demonstration trajectories
Learning Multi-Agent Hierarchical Policies in Unification with Subtask Discovery												
HSD; Yang et al [102]	decentralized agents with centralized learning	by maximizing Mutual Information (MI) between subtasks and trajectories	global	learned using r_ω	PRE, <i>sync</i>	n.r.	MI + task reward	D/C, S/L	D	HOM	discovery of diverse skills unified with the learning of hierarchical policies of all agents	the number of skills needs to be predefined and fixed
DOC; Chakravorty et al [98]	centralized agents	using policy gradients derived using the task reward	global	learned using task rewards	learned, <i>async</i>	n.r.	n.r.	D/C, S/L	D	HOM/HETR	learning Options hierarchy in a unified manner for all agents	the number of Options needs to be predefined and fixed // may perform poorly if the task reward is sparse

w.r.t. = *with respect to*; n.r. = *not required*; *sync* = synchronous, *async* = asynchronous.

D = discrete, C = continuous, S = small, L = large. PRE = predefined, DISC = discovered, HOM = homogeneous, HETR = heterogeneous.

All the task domains are assumed to have moderate to extreme reward sparsity.

Pateria et al [94] proposed an approach based on Options [11] to address both the above-mentioned issues. This approach is called Inter Subtask Empowerment based Multi-agent Options, or ISEMO. It exploits the inter-dependencies among heterogeneous agents to derive internal rewards called Inter Subtask Empowerment Rewards (ISER). The inter-dependencies are based on the predefined initiation conditions of various subtasks across different agents. ISER is given as a positive reward to an agent that *empowers* another agent by creating the environment conditions which satisfy the initiation conditions of the latter agent's subtask(s). ISER is added to the (external) task reward observed by the agent, thereby augmenting the task reward. Hence, even if the task reward is sparse for a particular agent, it can still learn useful behaviour using ISER for enabling other agents. Moreover, ISEMO uses the gradient-based learning of the termination function of each subtask on the lines of the Option-Critic [56]. The learned terminations are *asynchronous*. The authors compare ISEMO against the Cooperative HRL approach on a Search & Rescue task with a mix of heterogeneous and homogeneous agents. ISEMO is found to perform better than Cooperative HRL due to both ISER and the learned terminations. A similarity between ISEMO and Cooperative HRL is that ISEMO also uses the *decentralized agents with centralized learning* paradigm in which

the agents share the information about their local states and chosen subtasks. The main limitation of ISEMO is that the initiation conditions of various subtasks need to be manually defined for generating ISER. Also, ISEMO only learns the termination functions of the subtasks but requires the policies of the subtasks to be predefined.

Tang et al. [84] proposed a *deep learning* approach for MAHRL with predefined subtasks (defined using handcrafted subtask rewards). They use *decentralized agents with centralized learning* by adapting deep MARL approaches such as QMIX [78] and CommNet [55] for MAHRL. Different schemes for *synchronous and asynchronous* terminations of different subtasks are also proposed. This approach is evaluated on a modified version of the Trash collection task [25] and an online mobile game. It is found to outperform standard MARL and independently learned HRL agents. However, it is only shown to work for homogeneous agents.

The approaches discussed above are applicable to decentralized agents with centralized learning. Another approach is to learn *centralized agents* which have a common central policy at the highest-level. Ahilan et al [86] proposed a deep MAHRL approach in which a central manager policy chooses subtasks for a set of workers, simultaneously. The subtasks are predefined using subgoals. It is a feudal hierarchy in which each worker policy is a universal policy. Each worker policy belongs to a separate agent. Only the central manager receives the task rewards, while the workers learn using the subgoal achievement rewards. This approach is called Feudal Multi-agent Hierarchies (FMH). The workers can be heterogeneous or homogeneous. The terminations of different subtasks being performed by different workers are *synchronized* by the manager. FMH is evaluated on cooperative communication and navigation tasks in Multi-agent Particle Environments¹³ and found to outperform standard MARL.

The approaches discussed so far assume that the information can be shared among different agents either via communication or by using central policy. A more challenging but realistic paradigm is learning to coordinate *fully decentralized* agents. This paradigm does not assume the availability of communication or centralization during training or execution. Hence, the agents need to learn under *partial observability* [8]. Omidshafiei et al [42] introduced a theoretical framework for decentralized multi-agent coordination using subtasks (macro-actions) called Decentralized Partially Observable Semi-Markov MDPs (Dec-POSMDPs). Liu et al [53] provide an MAHRL approach based on Dec-POSMDPs. In this approach, the subtasks for all the agents are predefined and it is assumed that the agents have obtained a set of high-level trajectories from expert demonstrations containing sequences of subtasks. Then, the policy hierarchies of the decentralized agents are learned using an algorithm called Policy-based Expectation Maximization (PoEM) developed by the authors. All the policies take *history of state transitions* as input to resolve partial observability. PoEM is evaluated on multi-agent box pushing task and found to outperform dynamic-programming approaches. It works for *homogeneous or heterogeneous* agents. The subtask terminations can be *asynchronous*. However, a key limitation of this approach is that it requires the demonstration trajectories to be given in advance.

3.5.2 Learning Multi-Agent Hierarchical Policies in Unification with Subtask Discovery. Yang et al [102] proposed an MAHRL approach with skill¹⁴ discovery based on variational inference [68] (subsection 3.3.2). This approach is called Hierarchical Skill Discovery (HSD). The paradigm is *decentralized agents with centralized learning* where the agents are *homogeneous*. HSD uses QMIX [78] for centralized learning. In HSD, each agent uses a decentralized feudal hierarchy consisting of two levels. It selects a one-hot encoded skill vector using its higher-level policy. The lower-level universal policy takes this skill vector as an input and generates trajectories. This universal policy

¹³<https://github.com/openai/multiagent-particle-envs>

¹⁴Defined in subsection 2.3.

is learned by maximizing the Mutual Information (MI) between various input skill vectors and the corresponding trajectories, resulting in the discovery of *diverse skills*. However, an agent must also align the skill discovery with the coordination with other agents. Hence, HSD adds the task reward to the MI objective of the lower-level universal policy so that the learned skills are not only diverse but also grounded in the multi-agent task objective. The higher-level policy of each agent is learned using the task reward only. The skills of different agents terminate in a *synchronized* manner after fixed time steps. HSD is evaluated on a soccer game in STS2 simulator¹⁵. It performs better than standard MARL and independently trained HRL agents. The authors also report few identifiable discovered skills such as moving for offense, moving for defense, taking shots etc. They also find that the skills discovered only with MI objective and without adding the task reward do not lead to good performance. This shows that the discovered skills *cannot be task-agnostic*. The main limitation of HSD is that the number of skills/subtasks has to be predefined and fixed. There is no technique provided to grow this number or learn a large space of skills.

Chakravorty et al [98] proposed an MAHRL approach which extends the single-agent Option-Critic framework [56] for multi-agent coordination. This approach is called Distributed Option-Critic, or DOC. They use a single central higher-level policy for all the agents (hence, *centralized agents* paradigm). The lower-level Option policies are decentralized. The agents can be heterogeneous or homogeneous. The agents also have a learned broadcast function to share the local information with the central higher-level policy. The higher-level policy, lower-level Option policies, and the Option termination functions are learned using policy gradients derived using the task reward *only*. The learned terminations are *asynchronous*. The authors evaluate DOC on simple grid-world tasks and find that it performs competitively against standard MARL.

3.6 Key Takeaways from the Survey of Approaches

Through the survey of the wide-ranging selection of HRL approaches in this section, we aim to provide important insights about the expansive progress and the current state of the HRL research. In this regard, a few key takeaways from the surveyed approaches are discussed henceforth.

There is no consolidated framework to learn hierarchical policies, different approaches should be used for different benefits. The two sub-classes of the approaches for learning hierarchical policies provide different but important benefits. The feudal hierarchy approaches use subgoals (or instructions) to represent various possible subtasks [5, 50, 89, 76]. Such approaches have the scope to scale-up the subtask space by either using a large number of subgoals or using a low-dimensional *continuous* subgoal space. The scale-up is also possible because of the use of a universal policy at each level, which means that a large number of subtasks do not require many separate policies to be learned and stored, unlike the policy trees.

On the other hand, the policy tree approaches [11, 12, 10] are not constrained to learn only subgoal-based (or instruction-based [89]) subtasks, unlike the feudal approaches. Different subtask policies can represent different types of behaviours such as subgoal-based, trajectory-based (e.g. 'move in a circle'), general skills (e.g. 'drive through traffic') etc. *To summarize, different approaches provide different benefits but there is no consolidated framework which covers all the aspects.*

A variety of principles may be used for subtask discovery. There is no singular criteria for the quality of discovered subtasks. HRL researchers have a variety of ideas about the principle used for subtask discovery - partitioning state-action transition graphs [17, 24], clustering [22], diversifying exploration [63, 82], variational inference [68] etc. These principles are mostly

¹⁵<https://github.com/electronicarts/SimpleTeamSportsSimulator>

based on the objectives which are different from the task reward maximization objective. Hence, the optimality of the discovered subtasks can not be theoretically guaranteed. The task reward may itself be used for subtask discovery as shown by a few of the UNI approaches [56, 67, 77]. But, even among these approaches, there are different ways to include the reward into the subtask discovery process, e.g. using policy gradients [56, 67], using sub-optimality objective [77] etc. Hence, *there is no standard principle for subtask discovery yet. The reader must gauge a particular approach based on the preferred characteristics.*

New trends may emerge along the direction of large-scale skill discovery. In the recent years, the reinforcement learning community has shown a growing interest in the continuous control and robotic manipulation problems [73, 48]. Following this interest, HRL research has also begun to produce approaches for the discovery of robotic skills [68, 79] and learning hierarchical policies for continuous control [85, 76, 91]. This trend might continue in the near future, specially concerning the large-scale acquisition of robotic continuous-control skills, such as 'jumping', 'running', 'moving objects', 'driving through traffic', 'parking' etc., in a diversity of task domains [95] and the adaptation of the UNI approaches for the end-to-end learning of a large collection of such skills. This trend is important because diverse robotic skills are necessary for various real-life applications of HRL such as in warehouse management, surgical robots, autonomous driving etc.

4 OPEN PROBLEMS FOR FUTURE RESEARCH

HRL is garnering more interest in the reinforcement learning community. However, there still are different open problems concerning the scalability, efficiency, and theoretical robustness of HRL. A few such open problems are discussed henceforth, with the aim of identifying broader directions for future research. This discussion excludes the incremental issues and the extensions of any specific approach.

Building a lifelong knowledge base of transferable skills. Humans learn diverse skills during their lifetimes which are retained as valuable know-hows and selectively applied in different scenarios. The same faculty is essential for realizing the *general* HRL agents which can continuously learn, retain, and reuse a broad range of skills over their lifetimes. Building a lifelong knowledge base of skills requires the integration of multiple operations such as skill discovery, skill transfer, and hierarchical learning using the relevant skills as subtasks to perform a new task. Currently, the necessary operations are provided by disparate approaches, specifically for unsupervised skill discovery [68, 75, 95], skill transfer [60, 66], skill composition [87], skill reuse [100], and selective initiation of skills while training the HRL agent [101]. *Further research and development is necessary to design an integrated HRL framework which holistically combines these operations for maintaining the lifelong knowledge bases of transferable skills.*

Leveraging high-level planning for improved learning and adaptation. Replacing the model-free higher-level policy of an HRL agent with model-based planning, using the learned subtask-to-subtask (or subgoal-to-subgoal) transition dynamics, can potentially accelerate the learning of the agent on the long-horizons tasks and the adaptation to new tasks with shared dynamics [88, 81]. Zahavy et al [103] propose an approach to integrate high-level planning with HRL by using a set of sub-rewards to decompose the task reward. It is assumed that the sub-rewards are predefined. One subtask policy is defined for collecting one sub-reward. The higher-level planning heuristic finds the optimal sequence of subtasks by solving a Travelling Salesman Problem [2]. Sohn et al [81] propose **Neural Subtask Graph Solver** (NSGS) for planning. This approach uses a subtask graph in which the subtasks, their preconditions, and their inter-dependencies are predefined. The

higher-level planner is a graph solver, which plans over the subtask graph. The graph can be used to generalize to new tasks in a zero-shot fashion. The common limitation of both the approaches outlined above is that the subtasks and/or their inter-dependencies need to be handcrafted. Eysenbach et al [88] propose an approach, called Search on the Replay Buffer (SoRB), which automatically extracts the subgoal states from the experience replay buffer¹⁶ and builds a graph over them. In this graph, each node is a state and the weight of a transition edge between a pair of nodes is equal to the distance between the nodes estimated using the learned Q-value function. The higher-level planning heuristic finds the shortest-path over the subgoals in the graph, to reach a goal. The path is traversed using the lower-level primitive-action policy trained via reinforcement learning. While SoRB shows performance improvements over model-free reinforcement learning, it requires the primitive-action policy to be trained before the higher-level graph can be constructed for planning. Hence, SoRB is not end-to-end unified and it only addresses the problem of fast adaptation on the new goal-directed tasks, after learning the primitive-action policy and the higher-level graph on the previous goal-directed tasks.

To summarize, there are various aspects to consider when integrating high-level planning with HRL. Automatic discovery of subtasks is paramount. *More investigation is required on learning the subtask inter-dependencies automatically, instead of handcrafting them. New approaches should also be developed for the end-to-end unification of higher-level planning and lower-level policy learning.*

Providing theoretical support for HRL. HRL delivers apparent practical (or empirical) benefits over standard RL in many long-horizon task domains [76, 91, 67, 56, 50]. In theory, however, there is no incentive for using a hierarchical policy in terms of Q-value maximization alone. Standard reinforcement learning, especially Q-learning [4], is theoretically guaranteed to converge to the desired maximum Q-value [16] using only a primitive action policy. Majority of the *empirical* performance benefits of HRL may simply be due to improved exploration [28, 93]. Nachum et al [93], however, also show that a sophisticated non-hierarchical exploration scheme paired with a standard RL agent may provide equivalent performance as an HRL agent. Therefore, *more theoretical research needs to be done in order to support the empirical benefits of HRL and to identify the problem domains in which HRL provides clear advantage in terms of optimal performance.*

5 CONCLUSION

This survey provides a panoramic overview of the research done so far in the field of Hierarchical Reinforcement Learning (HRL), from the classical approaches to the recent advances. A novel and general taxonomy is devised that organizes the approaches into five broad classes based on the key challenges they address, which are: (i) Learning Hierarchical Policy (LHP), (ii) Learning Hierarchical Policy in Unification with Subtask Discovery (UNI), (iii) Independent Subtask Discovery (ISD), (iv) Transfer Learning with HRL (TransferHRL), and (v) Multi-agent HRL (MAHRL). Further classification of the approaches is also provided on the basis of the specific methodologies. Through the survey, we find that significant advances have been made with respect to addressing the two main sub-problems of HRL - how to learn a hierarchical policy and how to automatically discover subtasks - as mentioned in the HRL problem statement (subsection 2.2.3). The field is now making strides in the directions of transfer learning and multi-agent learning using HRL, in which new trends may emerge, specifically concerning subtask discovery.

Despite this noteworthy progress, there is scope for further growth in HRL to encourage its wider acceptability as a scalable and robust paradigm. In this regard, we identify a set of important open problems in order to motivate the future research and advancements. These chosen problems

¹⁶A storage of several experience data samples, of the form (s_t, a_t, r_t, s_{t+1}) , used to train an agent.

concern: (i) lifelong skill discovery and utilization for scalable HRL, (ii) improving the data efficiency by leveraging high-level planning, and (iii) providing more theoretical guarantees of optimality of the HRL approaches.

ACKNOWLEDGMENTS

This research was supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-RP-2020-019) and the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier-1 grant (19-C220-SMU-023).

REFERENCES

- [1] Richard Bellman. 1954. The theory of dynamic programming. *Bull. Amer. Math. Soc.*, 60, 6, (November 1954), 503–515. Retrieved from <https://projecteuclid.org/euclid.bams/1183519147>.
- [2] Richard Bellman. 1962. Dynamic programming treatment of the travelling salesman problem. *J. ACM*, 9, 1, (January 1962), 61–63. ISSN: 0004-5411. DOI: [10.1145/321105.321111](https://doi.org/10.1145/321105.321111).
- [3] Leonard E. Baum, Ted Petrie, George Soules, and Norman Weiss. 1970. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *Ann. Math. Statist.*, 41, 1, (February 1970), 164–171. DOI: <https://doi.org/10.1214/aoms/1177697196>.
- [4] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning*, 8, 3-4, 279–292. DOI: <https://doi.org/10.1007/BF00992698>.
- [5] Peter Dayan and Geoffrey E Hinton. 1993. Feudal reinforcement learning. In *Advances in Neural Information Processing Systems*. Volume 5. Morgan-Kaufmann, 271–278. Retrieved from <https://dl.acm.org/doi/10.5555/645753.668239>.
- [6] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*, 9, 8, (November 1997), 1735–1780. ISSN: 0899-7667. DOI: [10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735).
- [7] Ming Tan. 1997. *Multi-Agent Reinforcement Learning: Independent vs. Cooperative Agents*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 487–494. ISBN: 1558604952. Retrieved from <https://dl.acm.org/doi/10.5555/284860.284934>.
- [8] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artif. Intell.*, 101, 1–2, (May 1998), 99–134. ISSN: 0004-3702. Retrieved from <https://dl.acm.org/doi/10.5555/1643275.1643301>.
- [9] Oded Maron and Tomás Lozano-Pérez. 1998. A framework for multiple-instance learning. In *Advances in Neural Information Processing Systems*. Volume 10. MIT Press, 570–576. Retrieved from <https://dl.acm.org/doi/10.5555/302528.302753>.
- [10] Ronald Parr and Stuart Russell. 1998. Reinforcement learning with hierarchies of machines. In *Proceedings of the 1997 Conference on Advances in Neural Information Processing Systems 10 (NIPS '97)*. MIT Press, Denver, Colorado, USA, 1043–1049. ISBN: 0262100762. Retrieved from <https://dl.acm.org/doi/10.5555/302528.302894>.
- [11] Richard S. Sutton, Doina Precup, and Satinder Singh. 1999. Between mdps and semi-mdps: a framework for temporal abstraction in reinforcement learning. *Artif. Intell.*, 112, 1–2, (August 1999), 181–211. ISSN: 0004-3702. DOI: [10.1016/S0004-3702\(99\)00052-1](https://doi.org/10.1016/S0004-3702(99)00052-1).
- [12] Thomas G. Dietterich. 2000. Hierarchical reinforcement learning with the maxq value function decomposition. *J. Artif. Int. Res.*, 13, 1, (November 2000), 227–303. ISSN: 1076-9757. Retrieved from <https://dl.acm.org/doi/10.5555/1622262.1622268>.
- [13] Fredrik Linåker. 2000. Time series segmentation using an adaptive resource allocating vector quantization network based on change detection. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks (IJCNN'00)-Volume 6 - Volume 6 (IJCNN '00)*. IEEE Computer Society, USA, 6323. ISBN: 0769506194. DOI: [10.1109/IJCNN.2000.859416](https://doi.org/10.1109/IJCNN.2000.859416).
- [14] Rajbala Makar, Sridhar Mahadevan, and Mohammad Ghavamzadeh. 2001. Hierarchical multi-agent reinforcement learning. In *Proceedings of the Fifth International Conference on Autonomous Agents (AGENTS '01)*. Association for Computing Machinery, Montreal, Quebec, Canada, 246–253. ISBN: 158113326X. DOI: [10.1145/375735.376302](https://doi.org/10.1145/375735.376302).
- [15] Amy McGovern and Andrew G. Barto. 2001. Automatic discovery of subgoals in reinforcement learning using diverse density. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML '01)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 361–368. ISBN: 1558607781. Retrieved from <https://dl.acm.org/doi/10.5555/645530.655681>.
- [16] Francisco S Melo. 2001. Convergence of q-learning: a simple proof. *Institute of Systems and Robotics, Tech. Rep.*, 1–4. Retrieved from <http://users.isr.ist.utl.pt/~mtjspaan/readingGroup/ProofQlearning.pdf>.

- [17] Ishai Menache, Shie Mannor, and Nahum Shimkin. 2002. Q-cut - dynamic discovery of sub-goals in reinforcement learning. In *Proceedings of the 13th European Conference on Machine Learning (ECML '02)*. Springer-Verlag, Berlin, Heidelberg, 295–306. ISBN: 3540440364. Retrieved from <https://dl.acm.org/doi/10.5555/645329.650060>.
- [18] Khashayar Rohanimanesh and Sridhar Mahadevan. 2002. Learning to take concurrent actions. In *Proceedings of the 15th International Conference on Neural Information Processing Systems (NIPS'02)*. MIT Press, Cambridge, MA, USA, 1651–1658. Retrieved from <https://dl.acm.org/doi/10.5555/2968618.2968823>.
- [19] Martin Stolle and Doina Precup. 2002. Learning options in reinforcement learning. In *International Symposium on abstraction, reformulation, and approximation*. Springer, 212–223. DOI: https://doi.org/10.1007/3-540-45622-8_16.
- [20] Andrew G Barto and Sridhar Mahadevan. 2003. Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems*, 13, 1-2, 41–77. DOI: <https://doi.org/10.1023/A:1025696116075>.
- [21] G Dantzig and Delbert Ray Fulkerson. 2003. On the max flow min cut theorem of networks. *Linear inequalities and related systems*, 38, 225–231.
- [22] Bram Bakker and Jürgen Schmidhuber. 2004. Hierarchical reinforcement learning with subpolicies specializing for learned subgoals. In *Proceedings of the IASTED International Conference on Neural Networks and Computational Intelligence, NCI 2004, February 23-25, 2004, Grindelwald, Switzerland*. IASTED/ACTA Press, 125–130. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.218.3904&rep=rep1&type=pdf>.
- [23] Claudio Turchetti. 2004. *Stochastic models of neural networks*. Volume 102. IOS Press. Available at <https://www.iospress.nl/book/stochastic-models-of-neural-networks-2/>.
- [24] Özgür Şimşek, Alicia P. Wolfe, and Andrew G. Barto. 2005. Identifying useful subgoals in reinforcement learning by local graph partitioning. In (ICML '05). Association for Computing Machinery, Bonn, Germany, 816–823. ISBN: 1595931805. DOI: [10.1145/1102351.1102454](https://doi.org/10.1145/1102351.1102454).
- [25] Mohammad Ghavamzadeh, Sridhar Mahadevan, and Rajbala Makar. 2006. Hierarchical multi-agent reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 13, 2, 197–229. DOI: <https://doi.org/10.1007/s10458-006-7035-4>.
- [26] George Konidaris and Andrew Barto. 2007. Building portable options: skill transfer in reinforcement learning. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI'07)*. Morgan Kaufmann Publishers Inc., Hyderabad, India, 895–900. Retrieved from <https://dl.acm.org/doi/10.5555/1625275.1625420>.
- [27] Sridhar Mahadevan and Mauro Maggioni. 2007. Proto-value functions: a laplacian framework for learning representation and control in markov decision processes. *J. Mach. Learn. Res.*, 8, (December 2007), 2169–2231. ISSN: 1532-4435. Retrieved from <https://dl.acm.org/doi/10.5555/1314498.1314570>.
- [28] Nicholas K. Jong, Todd Hester, and Peter Stone. 2008. The utility of temporal abstraction in reinforcement learning. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 1 (AAMAS '08)*. International Foundation for Autonomous Agents and Multiagent Systems, Estoril, Portugal, 299–306. ISBN: 9780981738109. Retrieved from <https://dl.acm.org/doi/10.5555/1402383.1402429>.
- [29] Özgür Şimşek and Andrew G. Barto. 2008. Skill characterization based on betweenness. In *Proceedings of the 21st International Conference on Neural Information Processing Systems (NIPS'08)*. Curran Associates Inc., Vancouver, British Columbia, Canada, 1497–1504. ISBN: 9781605609492. Retrieved from <https://dl.acm.org/doi/10.5555/2981780.2981967>.
- [30] George Konidaris and Andrew Barto. 2009. Skill discovery in continuous reinforcement learning domains using skill chaining. In *Proceedings of the 22nd International Conference on Neural Information Processing Systems (NIPS'09)*. Curran Associates Inc., Vancouver, British Columbia, Canada, 1015–1023. ISBN: 9781615679119. Retrieved from <https://dl.acm.org/doi/10.5555/2984093.2984208>.
- [31] Stuart Russell and Peter Norvig. 2009. *Artificial Intelligence: A Modern Approach*. (3rd edition). Prentice Hall Press, USA. ISBN: 0136042597. Available at <https://dl.acm.org/doi/book/10.5555/1671238>.
- [32] Melike Baykal-Gürsoy. 2010. Semi-markov decision processes. *Wiley Encyclopedia of Operations Research and Management Science*. DOI: <https://doi.org/10.1002/9780470400531.eorms0757>.
- [33] Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. 2010. Why does unsupervised pre-training help deep learning? *J. Mach. Learn. Res.*, 11, (March 2010), 625–660. ISSN: 1532-4435. Retrieved from <https://dl.acm.org/doi/10.5555/1756006.1756025>.
- [34] Bernhard Hengst. 2010. Hierarchical reinforcement learning. In *Encyclopedia of Machine Learning*. Springer US, Boston, MA, 495–502. ISBN: 978-0-387-30164-8. DOI: https://doi.org/10.1007/978-0-387-30164-8_363.
- [35] Kfir Y. Levy and Nahum Shimkin. 2011. Unified inter and intra options learning using policy gradient methods. In *Proceedings of the 9th European Conference on Recent Advances in Reinforcement Learning (EWRL'11)*. Springer-Verlag, Athens, Greece, 153–164. ISBN: 9783642299452. DOI: [10.1007/978-3-642-29946-9_17](https://doi.org/10.1007/978-3-642-29946-9_17).
- [36] Alessandro Lazaric. 2012. Transfer in reinforcement learning: a framework and a survey. In *Reinforcement Learning*. Springer, 143–173. DOI: https://doi.org/10.1007/978-3-642-27645-3_5.

- [37] E. Todorov, T. Erez, and Y. Tassa. 2012. Mujoco: a physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033. doi: [10.1109/IROS.2012.6386109](https://doi.org/10.1109/IROS.2012.6386109).
- [38] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: encoder–decoder approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. Association for Computational Linguistics, Doha, Qatar, (October 2014), 103–111. doi: [10.3115/v1/W14-4012](https://doi.org/10.3115/v1/W14-4012).
- [39] Diederik P. Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings*. Retrieved from <http://arxiv.org/abs/1312.6114>.
- [40] Mostafa Al-Emran. 2015. Hierarchical reinforcement learning: a survey. *International journal of computing and digital systems*, 4, 02. doi: [http://dx.doi.org/10.12785/IJCDs/040207](https://doi.org/10.12785/IJCDs/040207).
- [41] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fiedjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *Nature*, 518, 7540, 529–533. doi: <https://doi.org/10.1038/nature14236>.
- [42] S. Omidshafiei, A. Agha-mohammadi, C. Amato, and J. P. How. 2015. Decentralized control of partially observable markov decision processes using belief space macro-actions. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 5962–5969. doi: [10.1109/ICRA.2015.7140035](https://doi.org/10.1109/ICRA.2015.7140035).
- [43] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning (Proceedings of Machine Learning Research)*. Volume 37. PMLR, Lille, France, (July 2015), 1889–1897. Retrieved from <http://proceedings.mlr.press/v37/schulman15.html>.
- [44] Marc G. Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Rémi Munos. 2016. Unifying count-based exploration and intrinsic motivation. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS'16)*. Curran Associates Inc., Barcelona, Spain, 1479–1487. ISBN: 9781510838819. Retrieved from <https://dl.acm.org/doi/10.5555/3157096.3157262>.
- [45] Christian Daniel, Herke Van Hoof, Jan Peters, and Gerhard Neumann. 2016. Probabilistic inference for determining options in reinforcement learning. *Machine Learning*, 104, 2-3, 337–357. doi: <https://doi.org/10.1007/s10994-016-5580-x>.
- [46] Jakob N. Foerster, Yannis M. Assael, Nando de Freitas, and Shimon Whiteson. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS'16)*. Curran Associates Inc., Barcelona, Spain, 2145–2153. ISBN: 9781510838819. Retrieved from <https://dl.acm.org/doi/10.5555/3157096.3157336>.
- [47] Karol Gregor, Danilo Jimenez Rezende, and Daan Wierstra. 2016. Variational intrinsic control. arXiv: [1611.07507](https://arxiv.org/abs/1611.07507). Retrieved from <http://arxiv.org/abs/1611.07507>.
- [48] Shixiang Gu, Ethan Holly, Timothy P. Lillicrap, and Sergey Levine. 2016. Deep reinforcement learning for robotic manipulation. *CoRR*, abs/1610.00633. arXiv: [1610.00633](https://arxiv.org/abs/1610.00633). Retrieved from <http://arxiv.org/abs/1610.00633>.
- [49] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross B. Girshick. 2016. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. arXiv: [1612.06890](https://arxiv.org/abs/1612.06890). Retrieved from <http://arxiv.org/abs/1612.06890>.
- [50] Tejas D. Kulkarni, Karthik R. Narasimhan, Ardavan Saeedi, and Joshua B. Tenenbaum. 2016. Hierarchical deep reinforcement learning: integrating temporal abstraction and intrinsic motivation. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS'16)*. Curran Associates Inc., Barcelona, Spain, 3682–3690. ISBN: 9781510838819. Retrieved from <https://dl.acm.org/doi/10.5555/3157382.3157509>.
- [51] Tejas D. Kulkarni, Ardavan Saeedi, Simanta Gautam, and Samuel J. Gershman. 2016. Deep successor reinforcement learning. arXiv: [1606.02396](https://arxiv.org/abs/1606.02396). Retrieved from <http://arxiv.org/abs/1606.02396>.
- [52] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2016. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings*. Retrieved from <https://arxiv.org/abs/1509.02971>.
- [53] Miao Liu, Christopher Amato, Emily P. Anesta, J. Daniel Griffith, and Jonathan P. How. 2016. Learning for decentralized control of multiagent systems in large, partially-observable stochastic environments. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*. AAAI Press, Phoenix, Arizona, 2523–2529. <https://dl.acm.org/doi/10.5555/3016100.3016253>.
- [54] Andrei A. Rusu, Sergio Gomez Colmenarejo, Çağlar Gülçehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. 2016. Policy distillation. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings*. Retrieved from <https://arxiv.org/abs/1511.06295>.

- [55] Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. 2016. Learning multiagent communication with backpropagation. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS'16)*. Curran Associates Inc., Barcelona, Spain, 2252–2260. ISBN: 9781510838819. Retrieved from <https://dl.acm.org/doi/10.5555/3157096.3157348>.
- [56] Pierre-Luc Bacon, Jean Harb, and Doina Precup. 2017. The option-critic architecture. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI'17)*. AAAI Press, San Francisco, California, USA, 1726–1734. Retrieved from <https://dl.acm.org/doi/10.5555/3298483.3298491>.
- [57] Carlos Florensa, Yan Duan, and Pieter Abbeel. 2017. Stochastic neural networks for hierarchical reinforcement learning. arXiv: [1704.03012](https://arxiv.org/abs/1704.03012). Retrieved from <http://arxiv.org/abs/1704.03012>.
- [58] Jakob N. Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2017. Counterfactual multi-agent policy gradients. arXiv: [1705.08926](https://arxiv.org/abs/1705.08926). Retrieved from <http://arxiv.org/abs/1705.08926>.
- [59] Roy Fox, Sanjay Krishnan, Ion Stoica, and Ken Goldberg. 2017. Multi-level discovery of deep options. arXiv: [1703.08294](https://arxiv.org/abs/1703.08294). Retrieved from <http://arxiv.org/abs/1703.08294>.
- [60] Kevin Frans, Jonathan Ho, Xi Chen, Pieter Abbeel, and John Schulman. 2017. Meta learning shared hierarchies. arXiv: [1710.09767](https://arxiv.org/abs/1710.09767). Retrieved from <http://arxiv.org/abs/1710.09767>.
- [61] Martin Klissarov, Pierre-Luc Bacon, Jean Harb, and Doina Precup. 2017. Learnings options end-to-end for continuous action tasks. *CoRR*, abs/1712.00004. arXiv: [1712.00004](https://arxiv.org/abs/1712.00004). Retrieved from <http://arxiv.org/abs/1712.00004>.
- [62] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. Curran Associates Inc., Long Beach, California, USA, 6382–6393. ISBN: 9781510860964. Retrieved from <https://dl.acm.org/doi/10.5555/3295222.3295385>.
- [63] Marlos C. Machado, Marc G. Bellemare, and Michael Bowling. 2017. A laplacian framework for option discovery in reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70 (ICML'17)*. JMLR.org, Sydney, NSW, Australia, 2295–2304. Retrieved from <https://dl.acm.org/doi/10.5555/3305890.3305918>.
- [64] Deepak Pathak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. 2017. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70 (ICML'17)*. JMLR.org, Sydney, NSW, Australia, 2778–2787. Retrieved from <https://dl.acm.org/doi/10.5555/3305890.3305968>.
- [65] Sainbayar Sukhbaatar, Ilya Kostrikov, Arthur Szlam, and Rob Fergus. 2017. Intrinsic motivation and automatic curricula via asymmetric self-play. arXiv: [1703.05407](https://arxiv.org/abs/1703.05407). Retrieved from <http://arxiv.org/abs/1703.05407>.
- [66] Chen Tessler, Shahar Givony, Tom Zahavy, Daniel J. Mankowitz, and Shie Mannor. 2017. A deep hierarchical approach to lifelong learning in minecraft. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI'17)*. AAAI Press, San Francisco, California, USA, 1553–1561. Retrieved from <https://dl.acm.org/doi/10.5555/3298239.3298465>.
- [67] Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, and Koray Kavukcuoglu. 2017. Feudal networks for hierarchical reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70 (ICML'17)*. JMLR.org, Sydney, NSW, Australia, 3540–3549. Retrieved from <https://dl.acm.org/doi/10.5555/3305890.3306047>.
- [68] Joshua Achiam, Harrison Edwards, Dario Amodei, and Pieter Abbeel. 2018. Variational option discovery algorithms. arXiv: [1807.10299](https://arxiv.org/abs/1807.10299). Retrieved from <http://arxiv.org/abs/1807.10299>.
- [69] Mordechai Ben-Ari and Francesco Mondada. 2018. Finite state machines. In *Elements of Robotics*. Springer, 55–61. doi: https://doi.org/10.1007/978-3-319-62533-1_4.
- [70] Zhiyuan Chen and Bing Liu. 2018. *Lifelong machine learning*. Number 3. Volume 12. Morgan & Claypool Publishers, 1–207. doi: [10.2200/S00737ED1V01Y201610AIM033](https://doi.org/10.2200/S00737ED1V01Y201610AIM033).
- [71] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. 2018. Diversity is all you need: learning skills without a reward function. arXiv: [1802.06070](https://arxiv.org/abs/1802.06070). Retrieved from <http://arxiv.org/abs/1802.06070>.
- [72] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research)*. Volume 80. PMLR, Stockholmssmässan, Stockholm Sweden, (October 2018), 1861–1870. Retrieved from <http://proceedings.mlr.press/v80/haarnoja18b.html>.
- [73] Tuomas Haarnoja, Aurick Zhou, Sehoon Ha, Jie Tan, George Tucker, and Sergey Levine. 2018. Learning to walk via deep reinforcement learning. arXiv: [1812.11103](https://arxiv.org/abs/1812.11103). Retrieved from <http://arxiv.org/abs/1812.11103>.
- [74] Jean Harb, Pierre-Luc Bacon, Martin Klissarov, and Doina Precup. 2018. When waiting is not an option: learning options with a deliberation cost. In *Proceedings of the AAAI Conference on Artificial Intelligence* number 1. Volume 32. Retrieved from <https://ojs.aaai.org/index.php/AAAI/article/view/11831>.
- [75] Karol Hausman, Jost Tobias Springenberg, Ziyu Wang, Nicolas Heess, and Martin A. Riedmiller. 2018. Learning an embedding space for transferable robot skills. In *6th International Conference on Learning Representations, ICLR*

- 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. OpenReview.net. Retrieved from <https://openreview.net/forum?id=rk07ZXZRb>.
- [76] Ofir Nachum, Shixiang Gu, Honglak Lee, and Sergey Levine. 2018. Data-efficient hierarchical reinforcement learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS'18)*. Curran Associates Inc., Montréal, Canada, 3307–3317. Retrieved from <https://dl.acm.org/doi/10.5555/3327144.3327250>.
- [77] Ofir Nachum, Shixiang Gu, Honglak Lee, and Sergey Levine. 2018. Near-optimal representation learning for hierarchical reinforcement learning. arXiv: [1810.01257](https://arxiv.org/abs/1810.01257). Retrieved from <http://arxiv.org/abs/1810.01257>.
- [78] Tabish Rashid, Mikayel Samvelyan, Christian Schröder de Witt, Gregory Farquhar, Jakob N. Foerster, and Shimon Whiteson. 2018. QMIX: monotonic value function factorisation for deep multi-agent reinforcement learning. arXiv: [1803.11485](https://arxiv.org/abs/1803.11485). Retrieved from <http://arxiv.org/abs/1803.11485>.
- [79] John D. Co-Reyes, Yuxuan Liu, Abhishek Gupta, Benjamin Eysenbach, Pieter Abbeel, and Sergey Levine. 2018. Self-consistent trajectory autoencoder: hierarchical reinforcement learning with trajectory embeddings. arXiv: [1806.02813](https://arxiv.org/abs/1806.02813). Retrieved from <http://arxiv.org/abs/1806.02813>.
- [80] Matthew Riemer, Miao Liu, and Gerald Tesauro. 2018. Learning abstract options. In *Advances in Neural Information Processing Systems*. Volume 31. Curran Associates, Inc., 10424–10434. Retrieved from <https://dl.acm.org/doi/10.5555/3327546.3327704>.
- [81] Sungryull Sohn, Junhyuk Oh, and Honglak Lee. 2018. Hierarchical reinforcement learning for zero-shot generalization with subtask dependencies. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS'18)*. Curran Associates Inc., Montréal, Canada, 7156–7166. Retrieved from <https://dl.acm.org/doi/10.5555/3327757.3327818>.
- [82] Sainbayar Sukhbaatar, Emily Denton, Arthur Szlam, and Rob Fergus. 2018. Learning goal embeddings via self-play for hierarchical reinforcement learning. arXiv: [1811.09083](https://arxiv.org/abs/1811.09083). Retrieved from <http://arxiv.org/abs/1811.09083>.
- [83] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: an introduction; 2nd ed. Adaptive computation and machine learning*. The MIT Press, Cambridge, MA. Available at <https://mitpress.mit.edu/books/reinforcement-learning-second-edition>.
- [84] Hongyao Tang, Jianye Hao, Tangjie Lv, Yingfeng Chen, Zongzhang Zhang, Hangtian Jia, Chunxu Ren, Yan Zheng, Changjie Fan, and Li Wang. 2018. Hierarchical deep multiagent reinforcement learning. arXiv: [1809.09332](https://arxiv.org/abs/1809.09332). Retrieved from <http://arxiv.org/abs/1809.09332>.
- [85] Z. Yang, K. Merrick, L. Jin, and H. A. Abbass. 2018. Hierarchical deep reinforcement learning for continuous action control. *IEEE Transactions on Neural Networks and Learning Systems*, 29, 11, 5174–5184. doi: [10.1109/TNNLS.2018.2805379](https://doi.org/10.1109/TNNLS.2018.2805379).
- [86] Sanjeevan Ahilan and Peter Dayan. 2019. Feudal multi-agent hierarchies for cooperative reinforcement learning. arXiv: [1901.08492](https://arxiv.org/abs/1901.08492). Retrieved from <http://arxiv.org/abs/1901.08492>.
- [87] Andre Barreto, Diana Borsa, Shaobo Hou, Gheorghe Comanici, Eser Aygün, Philippe Hamel, Daniel Toyama, Jonathan Hunt, Shihab Mourad, David Silver, and Doina Precup. 2019. The option keyboard: combining skills in reinforcement learning. In *Advances in Neural Information Processing Systems*. Volume 32. Curran Associates, Inc., 13052–13062. Retrieved from <https://proceedings.neurips.cc/paper/2019/file/251c5ffdb6b2cc21c446c963c76cf214-Paper.pdf>.
- [88] Ben Eysenbach, Russ R Salakhutdinov, and Sergey Levine. 2019. Search on the replay buffer: bridging planning and reinforcement learning. In *Advances in Neural Information Processing Systems*. Volume 32. Curran Associates, Inc., 15246–15257. Retrieved from <https://papers.nips.cc/paper/2019/file/5c48ff18e0a47baaf81d8b8ea51ecc92-Paper.pdf>.
- [89] YiDing Jiang, Shixiang (Shane) Gu, Kevin P Murphy, and Chelsea Finn. 2019. Language as an abstraction for hierarchical deep reinforcement learning. In *Advances in Neural Information Processing Systems*. Volume 32. Curran Associates, Inc., 9419–9431. Retrieved from <https://proceedings.neurips.cc/paper/2019/file/0af787945872196b42c9f73ead2565c8-Paper.pdf>.
- [90] Khimya Khetarpal and Doina Precup. 2019. Learning options with interest functions. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 01, (July 2019), 9955–9956. doi: [10.1609/aaai.v33i01.33019955](https://doi.org/10.1609/aaai.v33i01.33019955).
- [91] Andrew Levy, George Dimitri Konidaris, Robert Platt Jr., and Kate Saenko. 2019. Learning multi-level hierarchies with hindsight. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net. Retrieved from <https://openreview.net/forum?id=rzyECoAcY7>.
- [92] Matheus R. F. Mendonça, Artur Ziviani, and André M. S. Barreto. 2019. Graph-based skill acquisition for reinforcement learning. *ACM Comput. Surv.*, 52, 1, Article 6, (February 2019), 26 pages. issn: 0360-0300. doi: [10.1145/3291045](https://doi.org/10.1145/3291045).
- [93] Ofir Nachum, Haoran Tang, Xingyu Lu, Shixiang Gu, Honglak Lee, and Sergey Levine. 2019. Why does hierarchy (sometimes) work so well in reinforcement learning? arXiv: [1909.10618](https://arxiv.org/abs/1909.10618). Retrieved from <http://arxiv.org/abs/1909.10618>.

- [94] Shubham Pateria, Budhitama Subagdja, and Ah-Hwee Tan. 2019. Multi-agent reinforcement learning in spatial domain tasks using inter subtask empowerment rewards. In *IEEE Symposium Series on Computational Intelligence, SSCI 2019, Xiamen, China, December 6-9, 2019*. IEEE, 86–93. doi: [10.1109/SSCI44817.2019.9002777](https://doi.org/10.1109/SSCI44817.2019.9002777).
- [95] Archit Sharma, Shixiang Gu, Sergey Levine, Vikash Kumar, and Karol Hausman. 2019. Dynamics-aware unsupervised discovery of skills. arXiv: [1907.01657](https://arxiv.org/abs/1907.01657). Retrieved from <http://arxiv.org/abs/1907.01657>.
- [96] Shangtong Zhang and Shimon Whiteson. 2019. Dac: the double actor-critic architecture for learning options. In *Advances in Neural Information Processing Systems*. Volume 32. Curran Associates, Inc., 2012–2022. Retrieved from <https://proceedings.neurips.cc/paper/2019/file/4f284803bd0966cc24fa8683a34afc6e-Paper.pdf>.
- [97] Akhil Bagaria and George Konidaris. 2020. Option discovery using deep skill chaining. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. Retrieved from <https://openreview.net/forum?id=B1gqipNYwH>.
- [98] Jhelum Chakravorty, Patrick Nadeem Ward, Julien Roy, Maxime Chevalier-Boisvert, Sumana Basu, Andrei Lupu, and Doina Precup. 2020. Option-critic in cooperative multi-agent systems. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '20)*. International Foundation for Autonomous Agents and Multiagent Systems, Auckland, New Zealand, 1792–1794. ISBN: 9781450375184. Retrieved from <https://dl.acm.org/doi/abs/10.5555/3398761.3398984>.
- [99] Abhishek Gupta, Vikash Kumar, Corey Lynch, Sergey Levine, and Karol Hausman. 2020. Relay policy learning: solving long-horizon tasks via imitation and reinforcement learning. In *Proceedings of the Conference on Robot Learning (Proceedings of Machine Learning Research)*. Volume 100. PMLR, (30 Oct–01 Nov 2020), 1025–1037. Retrieved from <http://proceedings.mlr.press/v100/gupta20a.html>.
- [100] Leonard Hasenclever, Fabio Pardo, Raia Hadsell, Nicolas Heess, and Josh Merel. 2020. CoMic: complementary task learning & mimicry for reusable skills. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research)*. Volume 119. PMLR, (13–18 Jul 2020), 4105–4115. Retrieved from <http://proceedings.mlr.press/v119/hasenclever20a.html>.
- [101] Khimya Khetarpal, Martin Klissarov, Maxime Chevalier-Boisvert, Pierre-Luc Bacon, and Doina Precup. 2020. Options of interest: temporal abstraction with interest functions. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 4444–4451. Retrieved from <https://aaai.org/ojs/index.php/AAAI/article/view/5871>.
- [102] Jiachen Yang, Igor Borovikov, and Hongyuan Zha. 2020. Hierarchical cooperative multi-agent reinforcement learning with skill discovery. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*. International Foundation for Autonomous Agents and Multiagent Systems, 1566–1574. Retrieved from <https://dl.acm.org/doi/abs/10.5555/3398761.3398941>.
- [103] Tom Zahavy, Avinatan Hasidim, Haim Kaplan, and Yishay Mansour. 2020. Planning in hierarchical reinforcement learning: guarantees for using local policies. In (Proceedings of Machine Learning Research). Volume 117. PMLR, San Diego, California, USA, (August 2020), 906–934. Retrieved from <http://proceedings.mlr.press/v117/zahavy20a.html>.