Hierarchical Reinforcement Learning With Timed Subgoals

Anonymous Author(s)
Affiliation
Address
email

Abstract

Hierarchical reinforcement learning (HRL) holds great potential for sample-efficient learning on challenging long-horizon tasks. In particular, letting a higher level assign subgoals to a lower level has been shown to enable fast learning on difficult problems. However, such subgoal-based methods have been designed with static reinforcement learning environments in mind and consequently struggle with dynamic elements beyond the immediate control of the agent even though they are ubiquitous in real-world problems. In this paper, we introduce Hierarchical reinforcement learning with Timed Subgoals (HiTS), an HRL algorithm that enables the agent to adapt its timing to a dynamic environment by not only specifying what goal state is to be reached but also *when*. We discuss how communicating with a lower level in terms of such *timed subgoals* results in a more stable learning problem for the higher level. Our experiments on a range of challenging dynamic reinforcement learning environments show that our method is capable of sample-efficient learning where an existing state-of-the-art subgoal-based HRL method fails to learn stable solutions.[1]

1 Introduction

Hierarchical reinforcement learning (HRL) has recently begun to live up to its promise of sample-efficient learning on difficult long-horizon tasks. The idea behind HRL is to break down a complex problem into a hierarchy of more tractable subtasks. A particularly successful approach to defining such a hierarchy is to let a high-level policy choose a subgoal which a low-level policy is then tasked with achieving [6]. Due to the resulting temporal abstraction such *subgoal-based* HRL methods have been shown to be able to learn demanding tasks with unprecedented efficiency [29, 22, 17].

In order to realize the full potential of HRL it is necessary to design algorithms that enable concurrent learning on all levels of the hierarchy. However, the fact that the behaviour of the lower level changes during training introduces a major difficulty. From the perspective of the higher level, the reinforcement learning environment and the policy on the lower level constitute an effective environment which determines what consequences its actions will have. During training, the learning progress on the lower level renders this effective environment non-stationary. If this issue is not addressed, the higher level will usually not start to learn efficiently before the lower level is fully converged. This situation is similar to a manager and a worker trying to solve a task together while the meaning of the vocabulary they use for communication is continuously changing. Clearly, a stable solution can then only be found once the worker reacts reliably to instructions. Hence, to enable true concurrent learning, all levels in a hierarchy should see transitions that look like they were generated by interacting with a stationary effective environment.

[1] Videos and code, including our algorithm and the dynamic environments we evaluated it on, can be found here.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

Existing algorithms partially mask the non-stationarity of the effective environment by replacing the subgoal chosen by the higher level appropriately in hindsight. Combined with the subtask of achieving or progressing towards the assignend subgoal as fast as possible, this approach was shown to enable fast learning on a range of challenging sparse-reward, long-horizon tasks [22, 17]. What these methods do not take into account, however, is that, if adaptive temporal abstraction is used, the higher level in the hierarchy is effectively interacting with a semi-Markov decision process (SMDP), i.e., transition times vary. If the objective of the lower level is to reach a subgoal as fast as possible, then the amount of time that elapses until it reaches a given subgoal and returns control to the higher level will decrease during training. Hence, the distribution of the transition times the higher level sees will shift to lower values which introduces an additional source of non-stationarity. When trying to quickly traverse a static environment, such as a maze, this shift is in line with the overall task and will contribute to the learning progress.

Yet, as soon as dynamic elements that are beyond the immediate control of the agent are present in the environment, the situation changes radically. Consider, for example, the task of returning a tennis ball to a specified point on the ground by hitting it with a racket. This clearly requires the agent to time its actions so as to intercept the ball trajectory with the racket while it has the right orientation and velocity. Even if the higher level found a sequence of subgoals (specifying the state of the racket) that brought about the right timing, this solution would stop working as soon as the lower level learns to reach them faster. This would require the higher level to choose a different and possibly longer sequence of subgoals, a process that would continue until the lower level was fully converged. Hence, exposing the higher level to a non-stationary distribution of transition times will lead to training instability and slow learning. As it is the rule rather than the exception for real-world environments to contain dynamic elements beyond the immediate control of the agent – think of humans collaborating with a robot or an autonomous car navigating traffic – this problem can be expected to hinder the application of HRL to real-world tasks.

In order to solve the non-stationarity issue in dynamic environments, we propose to let the higher level choose not only what subgoal is to be reached but also *when*. By emitting such *timed subgoals*, consisting of a desired subgoal and a time interval that is supposed to elapse before it is reached, the higher level has explicit control over the transition times of the SMDP it interacts with. This completely removes the non-stationarity of transition times and allows for stable concurrent learning in dynamic environments when combined with a technique for hiding the non-stationiarity of transitions in the subgoal space. It furthermore gives the higher level explicit control over the degree of temporal abstraction, giving it more direct access to the trade off between a small effective problem horizon and exercising tight control over the agent. Communicating with the lower level in terms of timed subgoals moreover enables the agent to explicitly learn precise timing already early in the training process. In some tasks such as playing tennis or catching a ball this ability is crucial. The explicit representation of timing greatly facilitates credit assignment in these cases. Hence, the use of timed subgoals extends the class of tasks that can be solved efficiently by subgoal-based HRL methods.

The main contribution of this work is to distill these insights into the formulation of a sample-efficient HRL algorithm based on timed subgoals (Section 3) which we term Hierarchical reinforcement learning with Timed Subgoals (HiTS). We moreover propose 3 benchmark tasks that exemplify the challenges introduced by dynamic elements in the environment and demonstrate that HiTS succeeds in solving them while prior methods fail to learn a stable solution (Section 4). In a theoretical analysis, we show that the use of timed subgoals in combination with hindsight action relabeling [17] removes the non-stationarity of the SMDP generating the data the higher level is trained on (Section 3).

2 Background

2.1 Reinforcement learning and subgoal-based hierarchical reinforcement learning

We assume the reinforcement learning problem is formulated as a Markov decision process (MDP) defined by a tuple $(\mathcal{S}, \mathcal{A}, p, r, \rho_0, \gamma)$ consisting of state space $\mathcal{S}$, action space $\mathcal{A}$, transition probabilities $p(s' \mid s, a)$, reward function $r(s, a)$, initial state distribution $\rho_0(s_0)$ and discount factor $\gamma \in [0, 1)$. At each time step, the agent is provided with a state $s_t \in \mathcal{S}$ and samples an action $a_t \in \mathcal{A}$ from the distribution $\pi(\cdot \mid s_t)$ defined by its policy π. The environment then provides feedback in the form of a reward $r(s_t, a_t) \in \mathbb{R}$ and transitions to the next state $s_{t+1} \in \mathcal{S}$. The objective is to find an optimal

policy π^* that maximizes the expected discounted return

$$J(\pi) = \mathbb{E}_{\pi, s_0 \sim \rho_0} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right] . \tag{1}$$

We consider the goal-based HRL approach to solving this problem where a lower level is tasked with achieving a subgoal provided by a higher level [6, 22, 17]. The lower level (index 0) acts directly on the environment by outputting a primitive action in $\mathcal{A}$, whereas the higher level (index 1) proposes subgoals for level 0. Hence, the action of level 1 is the subgoal for level 0, i.e. $g^0 := a_t^1$. Since level 0 performs a variable number of steps τ in the environment before "finishing" its subtask, level 1 effectively interacts with a semi-Markov decision process (SMDP) [14, 15]. Its transition probabilities $p(s_{t+\tau}, \tau \mid s_t, a_t^1)$ determine the distribution not only of the next state $s_{t+\tau}$ level 1 observes but also of the time τ that elapses beforehand.

Existing subgoal-based algorithms reward the lower level either for being close to the subgoal [22] or for progressing in a given direction in subgoal space [29] or penalize every action that does not lead to the immediate achievement of the subgoal [17]. We consider the latter case of a shortest path objective with a reward

$$r^0(s', g^0) = \begin{cases} 0 & \text{if } \phi(s', g^0) = 1, \\ -1 & \text{otherwise} \end{cases}$$

for transitioning to a new state s'. The function $\phi(s', g^0)$ evaluates to 1 only if s' is sufficiently close to the subgoal g^0 in some metric (e.g. Euclidean distance). The higher level is queried for a new subgoal either if the lower level achieved a state sufficiently close to the subgoal or if the lower level has used up a fixed budget of actions.

2.2 Non-stationarity of the induced SMDP and hindsight action relabeling

The higher level of the hierarchy interacts with an SMDP that is induced by the environment dynamics and the behaviour of the lower level. Changes to the low-level policy during training render this SMDP non-stationary which prevents the higher level from learning a stable solution. Applying the recently proposed hindsight action relabeling technique can alleviate this problem for static environments [17]. We show that for dynamic environments the non-stationarity reappears.

For reasons of sample-efficiency, data should be reused after a policy update instead of being discarded. We therefore consider training with an off-policy algorithm that uses a replay buffer. On level 1 transitions of the form (s, a^1, r, s') are being stored, containing the state s, the action a^1, the reward r and the state s' at the end of the subtask execution. For now, we assume that the action space of level 1 (which is equal to the goal space of level 0) is the full state space.

Due to the changing low-level policy π^0, the distribution of the next state s' and the reward r given a fixed state action pair (s, a^1) will change in the course of training. Hence, the SMDP level 1 interacts with is non-stationary and learning a stable high-level policy becomes feasible only after the lower level is fully converged. *Hindsight action relabeling*, as proposed in [17], addresses this problem by "pretending" that the low-level policy is already optimal, i.e., is able to reach assigned subgoals. This is achieved by replacing the high-level action (corresponding to the assigned subgoal) with the state the lower level actually achieved before storing transitions in the replay buffer. If the reward depends on the action, it has to be adapted as well:

$$(s, a^1, r, s') \Longrightarrow (s, \hat{a}^1 := s', r^1(s, \hat{a}^1), s') \qquad \text{(hindsight action relabeling).} \tag{2}$$

In Fig. 1a we illustrate this relabeling procedure.

Proposition 1 *Applying hindsight action relabeling (Eq. 2) at level 1 generates transitions that follow a stationary state-transition and reward distribution provided that the subgoal space is equal to the full state space.*

The proof is given in Suppl. A. The remaining caveat is that the higher level should learn to restrict itself to subgoals which are feasible for the lower level. By conducting *testing transitions* [17], which are exempt from hindsight action relabeling and penalize infeasible subgoals, an appropriate incentive for the higher level can be introduced. Note that this reintroduces a non-stationarity to the data in the replay buffer. For the remaining part of this section we do not consider testing transitions as they are not directly related to the issue we are concerned with.

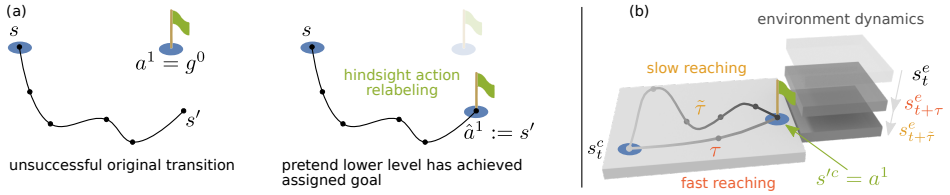

Figure 1: Relabeling results in a stationary state-transition distribution in static environments, but not in dynamic ones. (a) Action relabeling replaces the high-level action a^1 with the achieved state s' as if the lower level was already optimal. (b) In dynamic environments with a state that factorizes into a directly controllable part s^c and a remaining part s^e with dynamics, action relabeling does not prevent different environment outcomes for slow and fast low-level policies if the subgoal determines only s^c. In this example, the agent tries to reach a moving platform.

Dynamic Environments. We now consider environments with dynamic elements which are not directly controllable by the agent, for instance a flying ball or a moving platform. In this case, choosing the full state space as the subgoal space is problematic as it would include parts of the state the agent has no control over. Consequently, the higher level would be forced to learn the full dynamics of the environment in order to be able to propose feasible subgoals. In such a setting it is therefore more appropriate to restrict the subgoal space to the directly controllable part of the state. Formally, we assume the state space to factorize as $\mathcal{S} = \mathcal{S}^c \times \mathcal{S}^e$ into a directly controllable part $\mathcal{S}^c$, e.g. the agent, and a part corresponding to the rest of the environment $\mathcal{S}^e$ [16, 13, 9]. A typical choice for a subgoal space would then be $\mathcal{S}^c$, i.e. $\mathcal{G}^0 = \mathcal{S}^c$, as it allows the lower level to focus on controlling the agent alone and allows for transfer to related tasks. However, due to the restriction of the subgoal space, Prop. 1 does not apply anymore.

In particular, hindsight action relabling does not result in a stationary state-transition distribution anymore as illustrated in Fig. 1 (b). Even though relabeling the action of level 1 according to $a^1 = s'^c$ hides the influence of the lower level on the achieved goal, the state of the environment s'^e varies according to how long the lower level was in control. This transition time τ will shrink during the course of training as the lower level tries to optimize a shortest path objective. As a consequence, the same relabeled state-action pair $(s_t, \hat{a}^1 = s'^c)$ comes with different environmental outcomes $s^e_{t+\tau}$ and $s^e_{t+\tilde{\tau}}$ depending on the learning progress of the lower level. This makes it impossible for the agent to adapt its behavior to the environment before the lower level is fully converged. As a consequence, existing subgoal-based HRL methods with adaptive temporal abstraction struggle with dynamic environments.

3 Hierarchical reinforcement learning with timed subgoals

An observation that points to a way out of the non-stationarity dilemma in dynamic environments is that the part of the environment which is not under immediate control of the agent often only sparsely interacts with its readily controllable part. In the example from last section, the dynamics of the moving platform in Fig. 1 (b) is not influenced by the agent at all. As a result, knowing the elapsed time alone completely determines the distribution of s^e, no matter how complex its dynamics may be. We therefore propose to let time stand in for s^e and to condition the lower level on a *timed subgoal* $(g^0, \Delta t^0)$ where the desired time until achievement $\Delta t^0 \in \mathbb{N}_{>0}$ determines *when* the lower level is supposed to reach the subgoal state $g^0 \in S^c$.

Hence, after being assigned a timed subgoal at time t, the lower level stays in control until $s + \Delta t^0$. This definition of the hierarchy fixes the transition time the higher level sees to $\tau = \Delta t^0$. Combined with the assumption that s^e evolves independently of s^c and when replacing the action a^1 via hindsight action relabeling as discussed in Section 2.2, this completely removes the non-stationarity of the SMDP the level is interacting with.

Proposition 2 *If the not directly controllable part of the environment evolves completely independently from the controllable part, i.e., if $p\left(s'^e \mid s^e\right) = p(s'^e \mid s^e, s^c, a)$, and if hindsight action relabeling is used, the transitions in the replay buffer of a higher level assigning timed subgoals to a lower level are consistent with a completely stationary SMDP.*

The proof is given in Suppl. A. Of course, the assumption of an environment which is not at all influenced by the agent is restrictive. However, the interactions between agent and environment are often sparse in time (e.g. in the case of a racket hitting a ball). The higher level can then learn to identify the situations in which it has to exercise tight control over the agent because it is about to influence the environment and align its choice of desired subgoal achievement times to them. It then has full control over the interactions between agent and environment and the episode is divided into a sequence of time intervals during which the assumption of independent dynamics and consequently also Proposition 2 hold.

An existing subgoal-based method in conjunction with a state and goal space augmented with time, however, cannot provide sample-efficient learning on dynamic environments. Clearly, the objective of reaching a given time as fast as possible is vacuous. Furthermore, pursuing a timed subgoal which lies in the past will completely throw off the timing of the agent until the action budget of the lower level is exhausted. It is therefore necessary to design an HRL algorithm around the notion of timed subgoals to realize their potential for stable and efficient learning in dynamic environments.

3.1 The HiTS algorithm

In this section, we present *Hierarchical reinforcement learning with Timed Subgoals* (HiTS), an HRL algorithm for sample-efficient learning on dynamic environments. We consider a two-level hierarchy in which the higher level assigns timed subgoals to the lower level. For implementation details we refer to Suppl. B.

The objective of the higher level is to maximize the expected environment return while maintaining temporal abstraction. It outputs a timed subgoal $a^1 = (g^0, \Delta t_t^0) \in \mathcal{G}^0 \times \mathbb{N}_{>0}$ which the lower level then pursues for Δt_t^0 time steps. It then receives the cumulative environment reward plus a penalty $-c < 0$ for emitting a subgoal,

$$r^1(s_t, \Delta t_t^0) = \sum_{n=0}^{\Delta t_t^0 - 1} \gamma^n r(s_{t+n}, a_{t+n}^0) - c \,. \tag{3}$$

We additionally limit the subgoal budget of the higher level but bootstrap when this budget is exhausted. This ensures that even such failed episodes contribute to the learning progress by taking their final state into account.

The subtask assigned to the lower level is to achieve the timed subgoal $a^1 = (g^0, \Delta t_t^0)$. Note that the desired achievement time Δt_t^0 is given relative to the current time t, i.e., the goal state g^0 is to be achieved at the time $t + \Delta t_t^0$. It is therefore decremented in each time step, $\Delta t_{t+1}^0 = \Delta t_t^0 - 1$, before being passed to the policy. We chose this representation as the dynamics of an agent usually do not have an explicit time dependence and consequently time intervals rather than absolute times are relevant to control. In other words, a policy conditioned on this representation of a timed subgoal in combination with a time-invariant environment is automatically time-invariant as well.

The low-level policy sees the environment through the lens of an observation defined by a map $f^{\mathcal{O}} : \mathcal{S} \to \mathcal{O}^0$ (information hiding [6]). Furthermore, an environment state is translated into an achieved subgoal by a map $f^{\mathcal{G}} : \mathcal{S} \to \mathcal{G}^0$. In the setting of a factorized state space $\mathcal{S} = \mathcal{S}^c \times \mathcal{S}^e$ both maps may simply project onto the directly controllable part $\mathcal{S}^c$. The policy is conditioned on the observation and the desired timed subgoal and defines a density $\pi^0(\cdot \mid o_t^0, g^0, \Delta t_t^0)$ in action space. During execution of the hierarchy, an action $a_t^0 \in \mathcal{A}$ is sampled from this distribution and passed on to the environment. If the desired time until achievement has run out, i.e., if $\Delta t_{t+1}^0 = 0$, the higher level is queried for a new subgoal (see Fig. 2).

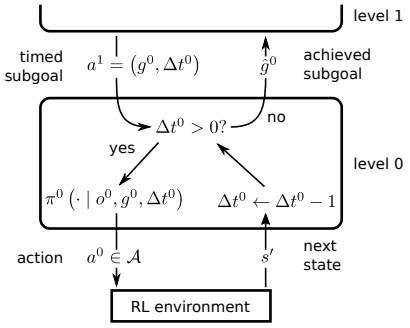

Figure 2: Control flow on the lower level of the HiTS hierarchy.

During training, the lower level receives a non-zero reward only if it achieves a timed subgoal, i.e., if the achieved subgoal is sufficiently close to the desired subgoal at the desired time of achievement. Hence, the reward in a transition to a new state s_{t+1} reads

$$r_t^0 = \begin{cases} 1 & \text{if } \phi(f^{\mathcal{G}}(s_{t+1}), g^0) = 1 \text{ and } \Delta t_{t+1}^0 = 0, \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

where $\phi : \mathcal{G}^2 \to [0, 1]$ specifies whether the achieved and the desired subgoal are sufficiently close to each other. Note that when learning the value function of the lower level we consider s_{t+1} as a terminal state when $\Delta t_{t+1}^0 = 0$ is reached, i.e., we do not bootstrap and set $\gamma^0 = 0$.

In order to realize the potential of HRL for sample-efficient learning, past experience should be reused instead of being discarded after each policy update. Hence, we use the off-policy reinforcement learning algorithm Soft Actor-Critic (SAC) [11] to train the policies on both levels. We address the non-stationarity issue by applying hindsight action relabeling to transitions before storing them in the replay buffer of the higher level,

$$(s_t, a_t^1 = (g_t^0, \Delta t_t^0), r_t, s_{t+\Delta t_t^0}) \Longrightarrow (s_t, \hat{a}_t^1 := (\hat{g}_t^0, \Delta t_t^0), r^1(s_t, \hat{a}_t^1), s_{t+\Delta t_t^0}) \ , \tag{5}$$

where $\hat{g}_t^0 = f^{\mathcal{G}}(s_{t+\Delta t_t^0})$ denotes the subgoal the lower level achieved. Note that the desired time until achievement Δt_t^0 does not have to be replaced as it is fixed by the higher level by definition. Under the assumption of a controllable part of the state which does not influence the rest of the environment, Proposition 2 ensures that the relabeled transitions in the replay buffer are consistent with a completely stationary SMDP.

In practice hindsight goal relabeling according to Eq. 5 is only applied if the desired subgoal was not achieved, i.e., if $\phi(\hat{g}^0, g^0) = 0$. This gives the higher level an idea of how precise the lower level is in achieving assigned subgoals while still hiding its failures. Furthermore, while pursuing a fixed percentage of timed subgoals, the lower level deterministically outputs the mean of its action distribution. If it still fails to achieve the desired goal, the higher level is penalized. Such testing transitions give the higher level an idea of the current capabilities of the low-level policy and therefore ensure that it assigns feasible subgoals [17]. While these two exceptions from hindsight action relabeling reintroduce a non-stationarity into the replay buffer of the higher level, they are necessary to make it respect the limitations of the lower level.

As the reward in Eq. 4 is sparse, we use *Hindsight Experience Replay* (HER) [1] to generalize from achieved to desired timed subgoals on the lower level. The resulting hindsight goal transitions are based on hindsight action transitions so as to conserve their stationarity property. When choosing an achieved subgoal at $n > 0$ time steps in the future as hindsight goal, the low-level transition is relabeled as

$$(s_t, a_t^0, r_t, s_{t+1}, g_t^0, \Delta t_t^0) \Longrightarrow (s_t, a_t^0, r_t, s_{t+1}, \hat{g}_t^0 := f^{\mathcal{G}}(s_{t+n}), \widehat{\Delta t}_t^0 := n) \ . \tag{6}$$

As the environments considered in the experiments are goal-conditioned, we also apply conventional HER on the higher level.

4 Experiments

The goal of our experimental evaluation is to compare HiTS with prior subgoal-based methods in terms of sample efficiency and stability of learning. We evaluate on four reinforcement learning environments[2] that exemplify the challenges introduced by dynamic elements which are beyond the immediate control of the agent (see Fig. 3). All environments require the agent to find the shortest path (in terms of time steps) to the goal, i.e., the reward is -1 in every time step unless the goal was achieved (in which case it is 0). If the agent did not reach the goal after a predefined number of time steps, the episode is terminated as well.

Platforms. The side-scroller-like *Platforms* environment requires the agent to learn to trigger the movement of a platform just at the right time to be able to use it later on to reach the goal. Hence, the timing of the agent's actions has a lasting impact on the dynamic elements in the environment which renders credit assignment in terms of primitive actions difficult. The Platforms environment is therefore – despite its simplicity – quite challenging for existing state-of-the-art methods both "flat" and hierarchical.

[2]The environments comply with OpenAI's gym interface and are available here.

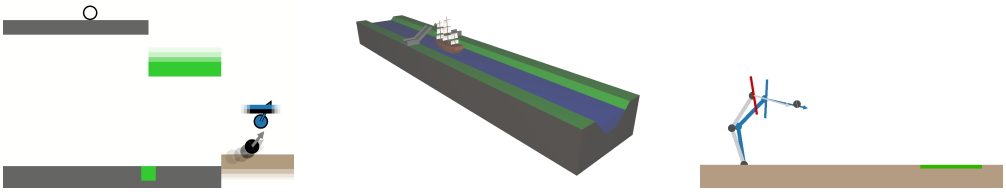

(a) Platforms (b) Drawbridge (c) Tennis2D

Figure 3: Dynamic environments. (a) To reach the goal, the agent has to trigger the movement of a platform at the right time. (b) In order not to collide with the drawbridge before it opens, the agent has to time the unfurling of its sails right. (c) A tennis ball is to be returned by a robot arm to a varying goal region.

Drawbridge. The agent in the Drawbridge environment has to unfurl the sails of a sailing ship at the right time to pass a drawbridge immediately after it opened. Note that the agent cannot actively decelerate the ship. It is therefore essential to not accelerate too early in order not to collide with the drawbridge and lose all momentum. While this control problem is quite simple, it requires the agent to wait – an ability existing subgoal-based HRL methods do not intrinsically have.

Tennis2D. In the *Tennis2D* environment a two-dimensional robot arm with a racket as an end effector has to return a ball to a specified area on the ground. As it is only through the contact between the racket and the ball that the agent can influence the outcome of the episode, it is crucial that the agent learns to precisely time the movement of the arm.

4.1 Empirical evaluation and comparison to baselines

We compare HiTS with two hierarchical and one "flat" baseline algorithm. As a subgoal-based HRL baseline we consider a two-level Hierarchical Actor-Critic (HAC) [17] hierarchy due to its capacity for concurrent learning and its sample efficiency. HAC uses a shortest path objective with the reward specified in Eq. 2.1 on both levels as well as hindsight action relabeling and HER. As a second baseline, we consider a two-level HAC hierarchy with a low-level observation and subgoal space that have been augmented with time. Thus, the subgoals used by this baseline hierarchy contain information about when they are to be achieved. The underlying HAC algorithm is left completely unchanged, however. Note that all hierarchies are given the same subgoal budget. We additionally consider SAC in combination with HER as a non-hierarchical baseline.

In order not to bias our comparison by the choice of underlying flat RL algorithm, we use our own implementation of HAC which is based on SAC. Hence, all considered hierarchies use the same flat RL algorithm. In Suppl. C we demonstrate that our implementation can match or exceed the performance of the original implementation of HAC on two static benchmark tasks introduced in [17]. Implementation details for HAC and HiTS can be found in Suppl. B. Details about training and hyperparameter optimization are given in Suppl. C.

Fig. 4a shows the average success rates of all policy hierarchies over the course of training for the Platforms environment. Note that all results reported in this section were obtained using deterministic policies that output the mean of the action distribution on each level. The success rates of the stochastic policies used during training can be found in Suppl. C. The SAC baseline fails to make any progress due to the lack of structured exploration and temporal abstraction. The two-level HAC hierarchy, on the other hand, begins to learn how to solve the task but stagnates at an averagge success rate of around 40%. Interestingly, augmenting the state, observation and goal spaces with absolute time does not improve the performance of the two-level HAC hierarchy but impedes it instead. As HAC is not aware of the significance of the time component, it continues to pursue a subgoal even if its time component already lies in the past. The solutions found by the augmented HAC hierarchy are therefore extremely brittle. Moreover, the use of absolute time requires the low-level policy to learn time-invariance from experience. In contrast to this, HiTS makes fast progress and reaches an asymptotic success rate of around 90%.

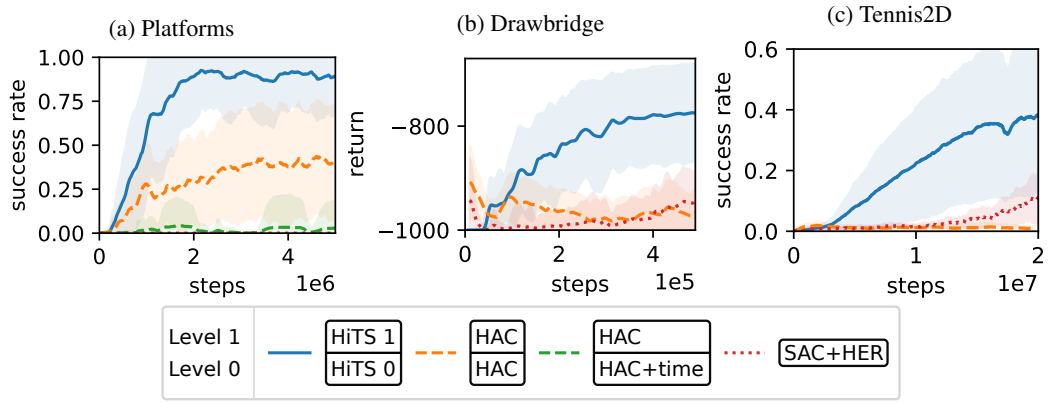

Figure 4: Results on the environments. We compare HiTS against HAC, HAC with time added to the goal space and SAC as a non-hierarchical baseline. Shaded regions indicate one standard deviation over results obtained with multiple random seeds.

A look at the learning progress of individual runs of the two-level HAC hierarchy (see Suppl. C) reveals that it does learn to solve the task in some cases. However, these solutions typically deteriorate quickly due to the learning progress on the lower level and the resulting non-stationarity of the SMDP as discussed in Section 2. In contrast to this, the performance of runs using the HiTS hierarchy is much more stable due to the use of timed subgoals as discussed in Section 3.

The challenge in the Drawbridge environment is to not only reach the goal (the end of the river) but to do so in the shortest possible time. Fig. 4b therefore shows the average return of all algorithms on the Drawbridge environment. In a successful episode it is equal to minus the time needed to reach the goal while an unsuccessful episode corresponds to a value of -1000 (the maximum episode length). As exploration is less challenging in this environment, the average return of the SAC policy consistently increases, albeit slowly. The performance of the two-level HAC hierarchy, on the other hand, decreases in expectation as it cannot reproduce the ideal behavior of waiting before unfurling the sails after the lower level has become fast at achieving subgoals. Increasing the subgoal budget would alleviate this problem but decrease temporal abstraction. Hence, a recursively optimal [3] HAC hierarchy may be far from optimal with respect to the environment MDP if the subgoal space is not equal to the full state space. HiTS does not suffer from this problem as the environment dynamics are independent from the agent's state in this case. Conditioning the lower level on time is therefore sufficient for HiTS to learn to time the passage through the drawbridge correctly.

In contrast to the environments discussed so far, controlling the agent in the Tennis2D environment is challenging as actions contain raw torques applied to the joints of the robot arm. Moreover, the trajectory as well as the spin of the ball varies considerably from episode to episode. Nevertheless, the HAC hierarchy initially manages to return a small fraction of the balls to the goal area (see Fig. 4c). Its performance decreases, however, after about 1.5 million time steps as the lower level has become too fast and a short sequence of subgoals cannot reproduce precisely timed movements anymore. As assigning credit for returning a ball to torques is challenging, SAC only learns slowly. The HiTS hierarchy can directly assign credit to a timed subgoal, i.e., to what robot state was to be achieved when. Fig. 5 shows how the discrepancy between the moment of contact between ball and racket and the closest achievement time of a timed subgoal decreases during training. The higher level can therefore exercise full control over the robot arm when it matters. As a consequence, HiTS achieves a performance of around 40% on average within 20 million time steps.

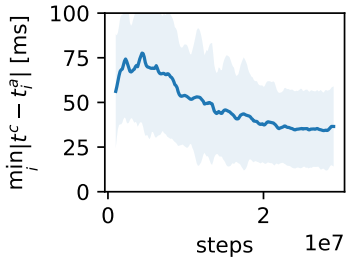

Figure 5: Absolute discrepancy between the time t^c of contact between ball and racket and the closest subgoal achievement time t_i^a in an episode of Tennis2D with HiTS.

The success rates and returns depicted in Fig. 4 exhibit a large variance over random seeds. This is caused in part by the environments in which a small change in the behavior of the agent can have a big impact on its performance. Additionally, all considered hierarchical algorithms suffer from

a failure mode caused by too aggressive exploration in the early stages of the training. This can result in the lower level not achieving any subgoals. As a result, the higher level does not learn to distinguish feasible from infeasible subgoals. This problem could be addressed either by conservative exploration or by tuning the entropy depending on the subgoal-achievement rate of the lower level.

5 Related work

The key role of abstraction in solving long-horizon, sparse-reward tasks has been recognized early on in the RL community [6, 27, 24, 28, 7]. The options framework [28, 25] formalizes the idea of augmenting the action space with closed-loop activities that take over control for an extended period of time. The resulting temporal abstraction facilitates credit assignment for a policy selecting among options. While options were originally handcrafted or learned separately [3, 25], the option-critic architecture [2] jointly learns the parameters of the options and a policy over options. However, to keep up stable temporal abstraction in this approach, regularization is necessary [2, 12]. A related approach to learning a hierarchy is based on the idea that a high-level controller should have a repertoire of skills to choose from [5] or to modulate [10, 13, 9, 8]. The set of skills is often pretrained using an auxiliary reward where diversity is encouraged by an information theoretic term [9, 8].

Alternatively, the lower level of the hierarchy may be trained to reach an assigned subgoal [26, 6, 25, 21]. In this setting, the lower level can be rewarded for moving in a specified direction in goal space [29], for being close to the goal state [22] or for achieving a goal state up to a tolerance [17]. In the latter case it is natural to query the higher level for a new subgoal only when the current one has been achieved or a fixed action budget is exhausted [17]. In contrast to this, using a dense reward is usually tied to a fixed temporal abstraction scheme [22, 29]. The goal space can either be learned [29, 23] or predefined using domain knowledge [22, 17].

It has been observed in the context of options [28, 25] as well as subgoal-based methods [22, 17] that the higher level in a hierarchy effectively interacts with a non-stationary SMDP. In order to enable off-policy learning from stored high-level transitions, Nachum et al. [22] approximately determine high-level actions that are likely to reproduce the stored behavior of the lower level under the current low-level policy. This procedure has to be repeated after each update to the low-level policy and introduces noise due to the use of sampling in the approximation. Levy et al. [17] address the non-stationarity issue by hindsight action relabeling as discussed in Section 2. None of these methods take into account, however, that, when using adaptive temporal abstraction as in [17], the transition times of the SMDP change over the course of training and consequently introduce an additional source of non-stationarity which prevents sample-efficient learning on dynamic environments. Although in Blaes et al. [4] the environment contains dynamic objects, they come to rest quickly if not manipulated by the agent, such that the non-stationarity issue does not need to be addressed.

In the context of robotics, via-points are used to specify a sequence of configurations a robot should move through, often also specifying when these configurations are to be reached [20]. They therefore serve a similar purpose as timed subgoals. The difference lies in how they are used. A classic approach in robotics would be to first obtain a controller for the robot, to then specify via-points, interpolate between them, and use the controller to follow the resulting trajectory. Hence, no concurrent learning is involved and the non-stationarity problem does not occur.

6 Conclusion

We present an HRL algorithm designed for sample-efficient learning in dynamic environments. In a theoretical analysis, we show how its use of timed subgoals in conjunction with hindsight action relabeling attenuates the non-stationarity problem of HRL even when the lower level is not conditioned on the full state. Moreover, our experiments on a range of challenging dynamic tasks demonstrate that our method outperforms existing state-of-the-art HRL baselines in terms of sample complexity and stability.

The use of timed subgoals in our algorithm enables the higher level to explicitly control the timing of the lower level. An evaluation of its potential for transfer of the low-level policy from one dynamic environment to another would therefore be an interesting direction for future work. Furthermore, using two or more levels which pursue timed subgoals may enable efficient learning on dynamic long-horizon tasks.

References

[1] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. P. Abbeel, and W. Zaremba. Hindsight experience replay. In *Advances in neural information processing systems*, pages 5048–5058, 2017.

[2] P.-L. Bacon, J. Harb, and D. Precup. The option-critic architecture. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

[3] A. G. Barto and S. Mahadevan. Recent advances in hierarchical reinforcement learning. *Discrete event dynamic systems*, 13(1):41–77, 2003.

[4] S. Blaes, M. V. Pogancic, J.-J. Zhu, and G. Martius. Control what you can: Intrinsically motivated task-planning agent. In *NeurIPS*, pages 12520–12531, 2019.

[5] C. Daniel, G. Neumann, and J. Peters. Hierarchical relative entropy policy search. In *Artificial Intelligence and Statistics*, pages 273–281. PMLR, 2012.

[6] P. Dayan and G. E. Hinton. Feudal reinforcement learning. In *Advances in neural information processing systems*, pages 271–278, 1993.

[7] T. G. Dietterich. Hierarchical reinforcement learning with the maxq value function decomposition. *Journal of artificial intelligence research*, 13:227–303, 2000.

[8] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine. Diversity is all you need: Learning skills without a reward function. In *International Conference on Learning Representations*, 2019.

[9] C. Florensa, Y. Duan, and P. Abbeel. Stochastic neural networks for hierarchical reinforcement learning. *arXiv preprint arXiv:1704.03012*, 2017.

[10] T. Haarnoja, K. Hartikainen, P. Abbeel, and S. Levine. Latent space policies for hierarchical reinforcement learning. In *International Conference on Machine Learning*, pages 1851–1860. PMLR, 2018.

[11] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.

[12] J. Harb, P.-L. Bacon, M. Klissarov, and D. Precup. When waiting is not an option: Learning options with a deliberation cost. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[13] N. Heess, G. Wayne, Y. Tassa, T. Lillicrap, M. Riedmiller, and D. Silver. Learning and transfer of modulated locomotor controllers. *arXiv preprint arXiv:1610.05182*, 2016.

[14] R. Howard. Semi-markovian decision processes. *Proc. Intern. Statist. Inst.*, page 625–652, 1963.

[15] Q. Hu and W. Yue. *Semi-Markov Decision Processes*, pages 105–120. Springer US, Boston, MA, 2008. doi: 10.1007/978-0-387-36951-8_5.

[16] G. Konidaris and A. Barto. Autonomous shaping: Knowledge transfer in reinforcement learning. In *Proceedings of the 23rd international conference on Machine learning*, pages 489–496, 2006.

[17] A. Levy, G. Konidaris, R. Platt, and K. Saenko. Learning Multi-Level Hierarchies with Hindsight. In *International Conference on Learning Representations*, 2019.

[18] A. Li, C. Florensa, I. Clavera, and P. Abbeel. Sub-policy adaptation for hierarchical reinforcement learning. In *International Conference on Learning Representations*, 2020.

[19] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

[20] K. M. Lynch and F. C. Park. *Modern Robotics*. Cambridge University Press, 2017.

[21] A. McGovern and A. G. Barto. Automatic discovery of subgoals in reinforcement learning using diverse density. In *Proceedings of the 18th International Conference on Machine Learning*, pages 361–368, 2001.

[22] O. Nachum, S. S. Gu, H. Lee, and S. Levine. Data-efficient hierarchical reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 3303–3313, 2018.

[23] O. Nachum, S. Gu, H. Lee, and S. Levine. Near-optimal representation learning for hierarchical reinforcement learning. In *International Conference on Learning Representations*, 2019.

[24] R. Parr and S. Russell. Reinforcement learning with hierarchies of machines. *Advances in neural information processing systems*, pages 1043–1049, 1998.

[25] D. Precup. *Temporal abstraction in reinforcement learning*. University of Massachusetts Amherst, 2000.

[26] J. Schmidhuber. Learning to generate sub-goals for action sequences. In *Artificial neural networks*, pages 967–972, 1991.

[27] A. Schwartz and S. Thrun. Finding structure in reinforcement learning. *Advances in neural information processing systems*, 7:385–392, 1995.

[28] R. S. Sutton, D. Precup, and S. Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.

[29] A. S. Vezhnevets, S. Osindero, T. Schaul, N. Heess, M. Jaderberg, D. Silver, and K. Kavukcuoglu. FeUdal networks for hierarchical reinforcement learning. volume 70 of *Proceedings of Machine Learning Research*, pages 3540–3549, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] Our claims are supported by the experiments in Section 4.
   (b) Did you describe the limitations of your work? [Yes] We discuss the remaining non-stationarity due to testing transitions and the subgoal achievement tolerance in Section 3.1. Furthermore, we describe a failure mode caused by too much exploration in Section 4.1.
   (c) Did you discuss any potential negative societal impacts of your work? [No] We do not see a possible immediate negative impact of our work.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] Our assumptions are discussed in Section 2 and Section 3.
   (b) Did you include complete proofs of all theoretical results? [Yes] Proofs can be found in Suppl. A.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] A link to a site with our code is given in a footnote in the abstract.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Training details are given in Suppl. C.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Fig. 4 contains shaded regions that correspond to one standard deviation computed over multiple runs with different random seeds. Details can be found in Suppl. C
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Information about used hardware and runtimes is given in Suppl. C.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] When presenting results on two static RL environments in Suppl. C we cite [17].
   (b) Did you mention the license of the assets? [Yes] We mention the license of the environments in Suppl. C.

(c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We include a link to our code of the algorithm and the environments in a footnote in the abstract.

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]