
Hierarchical Skills for Efficient Exploration

Jonas Gehring^{1,2} Gabriel Synnaeve¹ Andreas Krause² Nicolas Usunier¹

¹Facebook AI Research ²ETH Zürich

jgehring@fb.com

Abstract

In reinforcement learning, pre-trained low-level skills have the potential to greatly facilitate exploration. However, prior knowledge of the downstream task is required to strike the right balance between generality (fine-grained control) and specificity (faster learning) in skill design. In previous work on continuous control, the sensitivity of methods to this trade-off has not been addressed explicitly, as locomotion provides a suitable prior for navigation tasks, which have been of foremost interest. In this work, we analyze this trade-off for low-level policy pre-training with a new benchmark suite of diverse, sparse-reward tasks for bipedal robots. We alleviate the need for prior knowledge by proposing a hierarchical skill learning framework that acquires skills of varying complexity in an unsupervised manner. For utilization on downstream tasks, we present a three-layered hierarchical learning algorithm to automatically trade off between general and specific skills as required by the respective task. In our experiments, we show that our approach performs this trade-off effectively and achieves better results than current state-of-the-art methods for end-to-end hierarchical reinforcement learning and unsupervised skill discovery. Code and videos are available at <https://facebookresearch.github.io/hsd3>.

1 Introduction

A promising direction for improving the sample efficiency of reinforcement learning agents in complex environments is to pre-train low-level skills that are then used to structure the exploration in downstream tasks [23, 19, 27, 13, 30]. This has been studied in particular for the control of (simulated) robots, where there is a natural hierarchical decomposition of the downstream tasks into low-level control of the robot’s actuators with a skill policy, and a high-level control signal that specifies a direction or target robot configuration with coarser temporal resolution. The large body of work on unsupervised skill or option discovery in hierarchical reinforcement learning (HRL) for continuous control relies, explicitly or implicitly, on prior knowledge that low-level skills should control the center of mass of the robot [26, 17, 13, 43, 6]. This nicely fits a wide range of benchmark tasks that are variants of navigation problems, but the benefit of such hierarchical setups outside this problem class is unclear.

The prior knowledge embedded in a pre-trained skill defines a specific trade-off between sample efficiency and generality. Skills that severely constrain the high-level action space to elicit specific behavior (e.g., translation of the center of mass) are likely to provide the largest gains in sample efficiency, but are unlikely to be useful on a diverse set of downstream tasks. Conversely, low-level skills that expose many degrees of freedom are more widely applicable but less useful for guiding exploration. There is, thus, no single universally superior pre-trained skill. Depending on the downstream task, different skills might also be useful to efficiently explore in different parts of the environment.

In this paper, we aim to acquire skills that are useful for a variety of tasks while still providing strong exploration benefits. We propose to pre-train a *hierarchy of skills* of increasing complexity which can subsequently be composed with a high-level policy. In the context of simulated robots, each skill consists of controlling a part of the robot configuration over a short time horizon, such

as the position of the left foot of a humanoid, or the orientation of its torso. Skills of increasing complexity are constructed by jointly controlling larger portions of the configuration. These skills are modelled with a shared policy and pre-trained in an environment without rewards and containing the robot only. Subsequently, skills are used in downstream tasks within a three-level hierarchical policy: the highest level selects the skill (which specifies a *goal space*), the second level the target configuration within that skill (the *goal*), and the pre-trained skill performs the low-level control to reach the goal. Compared to standard approaches involving a single static pre-trained skill [13, 43], our approach offers increased flexibility for structuring exploration and offloads the issue of selecting prior knowledge from pre-training to downstream task learning. As a result, our skills can be acquired once per robot and applied to many different tasks.

We perform an experimental analysis of our hierarchical pre-training on a new set of challenging sparse-reward tasks with simulated bipedal robots. Our experiments show that each task is most efficiently explored by a distinct set of low-level skills, confirming that even on natural tasks, where locomotion is of primal importance, there is no overall single best pre-trained skill. We further show that dynamic selection of goal spaces with a three-level hierarchy performs equally or better than a generic skill on all tasks, and can further improve over the best single skills per task.

The main contributions of our work are summarized as follows:

- We propose a novel unsupervised pre-training approach that produces a hierarchy of skills based on control of variable feature sets.
- We demonstrate how to automatically select between different skills of varying complexity with a three-level hierarchical policy that selects skills, goals, and native actions.
- We introduce a benchmark suite of sparse-reward tasks that allows for consistent and thorough evaluation of motor skills and HRL methods beyond traditional navigation settings.
- We study the implications of prior knowledge in skills experimentally and showcase the efficacy of our hierarchical skill framework on the proposed benchmark tasks, achieving superior results compared to existing skill discovery and HRL approaches [13, 33, 53].

2 Related Work

The success of macro-operators and abstraction in classic planning systems [14, 40] has inspired a large body of works on hierarchical approaches to reinforcement learning [8, 46, 3, 50, 33, 39]. While the decomposition of control across multiple levels of abstraction provides intuitive benefits such as easier individual learning problems and long-term decision making, recent work found that a primary benefit of HRL, in particular for modern, neural-network based learning systems, stems from improved exploration capabilities [21, 35]. From a design perspective, HRL allows for separate acquisition of low-level policies (options; skills), which can dramatically accelerate learning on downstream tasks. A variety of works propose the discovery of such low-level primitives from random walks [26, 52], mutual information objectives [17, 13, 43, 6], datasets of agent or expert traces [37, 2, 25], motion capture data [36, 31, 42], or from dedicated pre-training tasks [15, 27].

In order to live up to their full potential, low-level skills must be useful across a large variety of downstream tasks. In practice, however, a trade-off between generality (wide applicability) and specificity (benefits for specific tasks) arises. A large portion of prior work on option discovery and HRL resolved this trade-off, explicitly or implicitly, in favor of specificity. This can be exemplified by the choice of test environments, which traditionally revolve around navigation in grid-world mazes [46, 8, 10]. In recent work on simulated robotics, navigation problems that are similar in spirit remain the benchmarks of choice [12, 33, 13, 27, 2]. In these settings, directed locomotion, i.e., translation of the robot’s center of mass, is the main required motor skill, and the resulting algorithms require a corresponding prior. This prior is made explicit with skills partitioning the state space according to the agent’s position [13, 43], or is implicitly realized by high contribution of position features to reward signals [33, 34, 2] (Appendix G). Similarly, works that target non-navigation environments acquire skills with tailored pre-training tasks [38] or in-domain motion capture data [36, 31, 32]. In contrast, our work is concerned with learning low-level skills without extra supervision from traces or pre-training task design, and which do not prescribe a fixed trade-off towards a particular type of behavior.

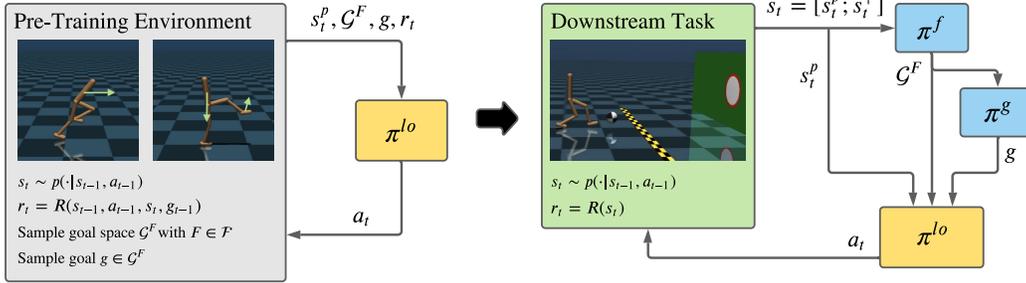


Figure 1: Illustration of our proposed hierarchical skill learning framework. **Left** Low-level policies are learned in an empty pre-training environment, with the objective to reach random configurations (goal g) of a sampled skill (goal space \mathcal{G}^F defined over a feature set F). Examples of goal space features are translation along the X-axis or the position of a limb. **Right** Learning on downstream tasks with a three-level hierarchical policy to select a goal space, a goal and finally a native action a_t with the pre-trained low-level policy. The low-level policy acts on proprioceptive states s^p , while high-level policies π^f and π^g leverage extra task-specific information via s^+ .

3 Hierarchical Skill Learning

3.1 Overview

In this work, we propose an approach where we first acquire low-level policies that are able to carry out a useful set of skills in an unsupervised manner, i.e., without reward signals from the main tasks of interest. We subsequently employ these skills in a hierarchical reinforcement learning setting (Figure 1). The pre-training environment consists solely of the robot. In downstream tasks, additional objects may be present besides the robot; these may be fixed (e.g., obstacles) or only indirectly controllable (such as a ball).

Crucially, the usefulness of a skill policy stems from providing additional structure for effective exploration. In order to provide benefits across a wide range of tasks, skills need to support both fast, directed exploration (e.g., locomotion) as well as precise movements (e.g., lift the left foot while bending down). We propose to fulfill these requirements with short-horizon, goal-directed low-level policies that are trained to achieve target configurations of robot-level state features such as the position or orientation of its torso or relative limb positions. We denote this feature space with \mathcal{S}^g . This allows the definition of a hierarchy of skills by controlling single features and their combinations, resulting in varying amounts of control exposed to high-level policies. Each skill is trained to reach goals in a goal space \mathcal{G}^F defined over a set of features F of \mathcal{S}^g , yielding policies of the form

$$\{\pi_F^{lo} : \mathcal{S}^p \times \mathcal{G}^F \rightarrow A\}_{F \in \mathcal{F}}$$

with proprioceptive observations \mathcal{S}^p (Figure 1, left). Goal-directed skill policies are trained without task-specific rewards, relying solely on state features related to the specific robot and the prescribed hierarchy of goal spaces.

On downstream tasks, high-level policies operate with a combinatorial action space as the skill policy π^{lo} is conditioned on both a feature set F and a concrete goal $g \in \mathcal{G}^F$. State spaces are enriched with task-specific features \mathcal{S}^+ , s.t. $\mathcal{S} = \mathcal{S}^p \cup \mathcal{S}^+$, that contain information regarding additional objects in the downstream task. This enriched state space is only available to the high-level policy. We model the high-level policy in a hierarchical fashion with two policies

$$\pi^f : \mathcal{S} \rightarrow \mathcal{F}, \quad \pi^g : \mathcal{S} \times \mathcal{F} \rightarrow \mathcal{G}^F,$$

that prescribe a goal space and a goal, respectively (Figure 1, right). With a task- and state-dependent policy π^f , it is possible to not only select the required set of skills for a given environment, but also to switch between them within an episode. In this three-level hierarchy, the higher-level policy π^f explores the hierarchy of goal spaces and dynamically trades off between generality and specificity. Temporal abstraction is obtained by selecting new high-level actions at regular intervals but with a reduced frequency compared to low-level actions. We term the resulting learning algorithm *HSD-3*, emphasizing the three different levels of control (goal spaces, goals and native actions) obtained after our hierarchical skill discovery phase.

3.2 Unsupervised Pre-Training

During pre-training, skill policies act in an MDP \mathcal{M} with proprioceptive states s^p , native actions a , transition probabilities and initial states $s_0^p \in \mathcal{S}_0^p$ drawn from a distribution P_0 [45, for example], but without an extrinsic reward function. The skill policies are trained to reach goals in a goal space defined over S^g consisting of robot-level sensor readings, which may include non-proprioceptive information.

Critically, each skill π_F^{lo} aims to achieve goals defined over a variable-sized subset of features F of S^g . To avoid learning a combinatorial number of independent policies, we follow the principle of universal value function approximators [41] and share the parameters of each skill. We augment the policy’s input with goal and goal space information, and learn a single set of parameters θ , leading to $\pi_\theta^{\text{lo}} : \mathcal{S}^p \times \mathcal{F} \times \mathcal{G} \rightarrow \mathcal{A}$, where $\mathcal{F} := \{F : F \subseteq \mathcal{S}^g\}$ and $\mathcal{G} := \{G^F : F \subseteq \mathcal{S}^g\}$. For ease of modeling, F is provided to π_θ^{lo} as a bag-of-words input where a coordinate is set to 1 if its respective feature is included in F and to 0 otherwise. Also, the low-level policy receives goals relative to the current values of the s^g . Thus, at each time step, the input goal to the policy is updated to reflect the progress towards the goal. As a learning signal, we provide a distance-based reward as $R(s^g, a, s^{g'}, F, g) := \|\omega^F(s^g) - g\|_2 - \|\omega^F(s^{g'}) - g\|_2$, where $\omega^F : \mathcal{S}^g \rightarrow \mathcal{G}^F$ is a fixed transformation that selects the subset F of goal space features and applies a suitable normalization. The selection of features for \mathcal{S}^g and transformations ω^F represents the prior knowledge that we utilize for unsupervised pre-training. We discuss examples for bipedal robots in our experiments and in Appendix B.

During training, features F and goals g are sampled anew for each episode in the environment. Episodes consist of several dozen steps only, reflecting the fact that we are interested in short-horizon skills. In-between episodes, we reset the simulation to states drawn from \mathcal{S}_0^p only sporadically to encourage the resulting policies to be applicable in a wide range of states, facilitating their sequencing under variable temporal abstractions. Likewise, rewards are propagated across episodes as long as the simulation is not reset. The low-level policy parameters θ can be optimized with any reinforcement learning algorithm; here, we opt for Soft Actor-Critic (SAC), a state-of-the-art method tailored to continuous action spaces [18]. Pseudo-code for the pre-training algorithm and further implementation details are provided in Appendix C.

3.3 Hierarchical Control

After a low-level policy π^{lo} has been obtained via unsupervised pre-training, we employ it to construct a full hierarchical policy $\pi(a|s) = \pi^f(F|s)\pi^g(g|s, F)\pi^{\text{lo}}(a|s^p, F, g)$. The hierarchical policy is applied to downstream tasks, where we are supplied with a reward signal and state observations $\mathcal{S} = \mathcal{S}^p \cup \mathcal{S}^+$. While different frequencies may be considered for acting with π^f and π^g , in this work both high-level actions are selected in lock-step. For jointly learning π^f and π^g , we extend the SAC algorithm to incorporate the factored action space $\mathcal{F} \times \mathcal{G}$. In particular, it is necessary to extend SAC to deal with a joint discrete and continuous action space, which we describe below (see Appendix D for further details and pseudo-code for high-level policy training).

We formulate our extension with a shared critic $Q(s, F, g)$, incorporating both factors of the action space. $Q(s, F, g)$ is trained to minimize the soft Bellman residual [18, Eq. 5], with the soft value function $V(s)$ including entropy terms for both policies. Since π^f outputs a distribution over discrete actions (to select from available goal spaces), we compute its entropy in closed form instead of using the log-likelihood of the sampled action as for π^g . We employ separate temperatures α and β^F for the goal space policy as well as for each corresponding goal policy; β^F are normalized to account for varying action dimensions across goal spaces.

$$V(s) = \sum_{F \in \mathcal{F}} \pi^f(F|s) \mathbb{E}_{g \sim \pi^g} \left[Q(s, F, g) - \frac{\beta^F}{|F|} \log \pi^g(g|s, F) \right] + \alpha (\mathcal{H}(\pi^f(\cdot|s)) - \log |\mathcal{F}|).$$

Compared to the usual SAC objective for continuous actions, we subtract $\log |\mathcal{F}|$ from the entropy to ensure a negative contribution to the reward. This matches the sign of the entropy penalty from π^g , which is negative in standard hyper-parameter settings¹ [18]. Otherwise, the entropy terms for discrete and continuous actions may cancel each other in Q-function targets computed from $V(s)$.

¹The entropy for continuous actions sampled from a Gaussian distribution can be negative since it is based on probability densities.

The policy loss [18, Eq. 7] is likewise modified to include both entropy terms. Parameterizing the policies with ϕ and ψ and the Q-function with ρ , we obtain

$$J_{\pi^{f,g}}(\phi, \psi) = \mathbb{E}_{s \sim B} \left[\sum_{F \in \mathcal{F}} \pi_{\phi}^f(F|s) \mathbb{E}_{g \sim \pi_{\psi}^g} \left[\frac{\beta^F}{|F|} \log \pi_{\psi}^g(g|F, s) - Q_{\rho}(s, F, g) \right] - \alpha \mathcal{H}(\pi_{\phi}^f(\cdot|s)) \right],$$

where the replay buffer is denoted with B . The scalar-valued entropy temperatures are updated automatically during learning to maintain target entropies $\bar{\mathcal{H}}^f$ and $\bar{\mathcal{H}}^g$ for both π^f and π^g . The following losses are minimized at each training step:

$$J(\alpha) = \mathbb{E}_{s \sim B} \left[\alpha \left(\mathcal{H}(\pi^f(\cdot|s)) - \bar{\mathcal{H}}^f \right) \right]$$

$$J(\beta) = \mathbb{E}_{s \sim B} \left[- \sum_{F \in \mathcal{F}} \beta^F \pi^f(F|s) \mathbb{E}_{g \sim \pi^g} \left[\frac{1}{|F|} \log \pi^g(g|s, F) + \bar{\mathcal{H}}^g \right] \right].$$

We implement temporal abstraction by selecting high-level actions with a reduced frequency. This incurs a reduction in available training data since, when taking a high-level action every c steps in an N -step episode, we obtain only N/c high-level transitions. To leverage all available transitions gathered during training, we adopt the step-conditioned critic proposed by Whitney et al. [52] in our SAC formulation. The Q-function receives an additional input $0 \leq i \leq c$ marking the number of steps from the last high-level action, and is trained to minimize a modified soft Bellman residual:

$$J_Q(\rho) = \mathbb{E}_{\substack{F_t, g_t, i, \\ s_t, \dots, t+c-i, \\ a_t, \dots, t+c-i \sim B}} \left[\frac{1}{2} \left(Q_{\rho}(s_t, F_t, g_t, i) - \left(\sum_{j=0}^{c-i-1} (\gamma^j r(s_{t+j}, a_{t+j})) + \gamma^{c-i} V(s_{t+c-i}) \right) \right)^2 \right]$$

As in [52], $Q(s, F, g, 0)$ is used when computing $V(s)$ and the policy loss $J_{\pi^{f,g}}(\phi, \psi)$.

The factorized high-level action space of HSD-3 has previously been studied as parameterized action spaces [28, 51], albeit with a small number of discrete actions and not in the context of hierarchical RL. A possible application of Soft Actor-Critic is described in Delalleau et al. [9]; our approach differs in that we (a) compute the soft value function as a weighted sum over all discrete actions, and (b) opt for two separately parameterized actors and a shared critic. These design choices proved to be more effective in initial experiments.

4 Benchmark Environments

We propose a benchmark suite for comprehensive evaluation of pre-trained motor skills, tailored to bipedal robots (Figure 2) and implemented for the MuJoCo physics simulator [49, 48]. We selected tasks that require diverse abilities such as jumping (Hurdles), torso control (Limbo), fine-grained foot control (Stairs, GoalWall) and body balance (PoleBalance). Besides PoleBalance, all tasks also require locomotion. The tasks are designed to have a sparse reward that is provided once an obstacle or step has been passed or a goal has been shot, emphasizing the need for good exploration. As an exception, PoleBalance provides a constant reward during the entire episode, which ends when the pole placed on the robot falls over. PoleBalance requires precise motor control, and is, in combination with its reward structure, chiefly outside of the realm of benchmarks that have been traditionally used to evaluate HRL techniques. Obstacle position and step lengths are subject to random perturbations, sampled anew for each episode. We tailor several environment parameters to the concrete robot being controlled, e.g., the positions of Limbo bars. For the effectively two-dimensional Walker robot, we limit the movement of objects (ball, pole) to translation along X and Z and rotation around the Y-axis. In all environments, we label simulation states that would cause the robot to fall over as invalid; these will terminate the episode with a reward of -1. Appendix A contains full specifications for all tasks.

Variants of some of our benchmark tasks have been proposed in prior works: Hurdles [20, 24, 13, 38], Stairs [29, 16], Gaps [20, 48], and PoleBalance is inspired by the classical CartPole problem [4]. Our tasks are different from earlier published counterparts in parameterization and in the sparsity of rewards they provide. All environments are provided via a standard Gym interface [5] with accompanying open-source code, enabling easy use and re-use.

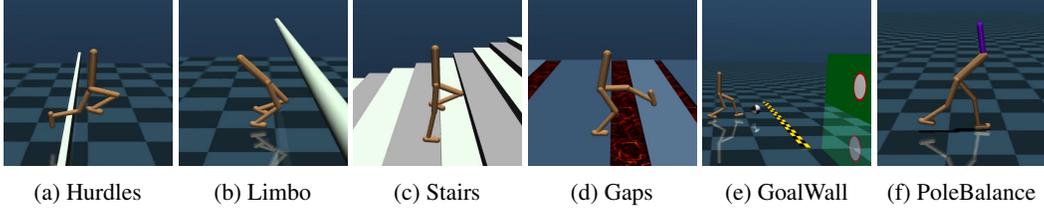


Figure 2: Benchmark environments for evaluating motor skills for bipedal robots, pictured with the Walker robot. Hurdle and Limbo bar heights and spacing, as well as stair lengths, are sampled randomly from fixed distributions. We also include a combination of Hurdles and Limbo (HurdlesLimbo), in which both obstacle types alternate.

5 Experimental Results

For our experiments, we first pre-train skill policies as described in Section 3.2 in an empty environment. The feature set to describe target configurations consists of the robot’s translation along the X-axis, its position on the Z-axis, torso rotation around the Y-axis and positions of each foot, relative to the body (Appendix B). With these five features, we obtain a set of $2^5 - 1 = 31$ skills with a shared set of policy parameters. We then train separate high-level policies for each benchmark environment, directing the same set of pre-trained skills. High-level actions are taken every 5 environment steps; in the PoleBalance environment, which requires short reaction times, all three policies operate at the same frequency. All neural network policies employ skip connections as proposed for SAC by Sinha et al. [44]. We show results over 9 seeds for each run; in experiments involving pre-training, we combine 3 seeds for pre-training with 3 seeds for high-level policy training. Pre-training takes approximately 3 days on 2 GPUs (V100) and fine-tuning (downstream task training) takes 2 days to reach 5M samples on 1 GPU. Further details regarding the training setup, hyper-parameters for skill and high-level policy training, as well as for baselines, can be found in Appendix E.

5.1 Trade-offs for Low-Level Skills

In a first experiment with the Walker robot, we aim to investigate the trade-off between specificity and generality of low-level skills. For each environment, we train high-level policies directing a *fixed* skill within its respective goal space, which for example consists of torso-related features only, or of feet positions combined with forward translation. At the end of training, we measure average returns obtained for all 31 skills in a deterministic evaluation setting, and plot quartiles over different seeds for the best skills in Figure 3. The performance for different fixed skills varies significantly across environments, and *no single skill is able to obtain the best result (shaded) across all environments*. While controlling X, Y and Z features only produces the good hierarchical policies for most tasks, poor results are obtained in Gaps and GoalWall. Controlling feet positions (LF and RF) is crucial for GoalWall, although close-to-optimal performance is only achieved in few runs across the respective

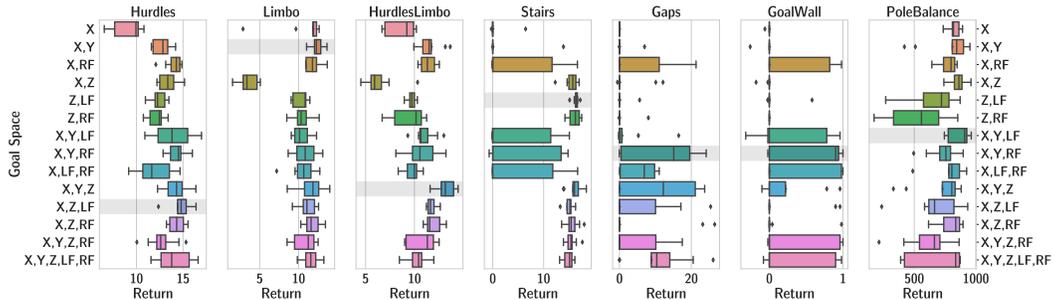


Figure 3: Returns achieved after 5M samples on the benchmark tasks with the Walker robot with fixed low-level policy goal spaces (quartiles). Each row corresponds to a set of features for the respective goal space. Best skills (marked) differ significantly across tasks.

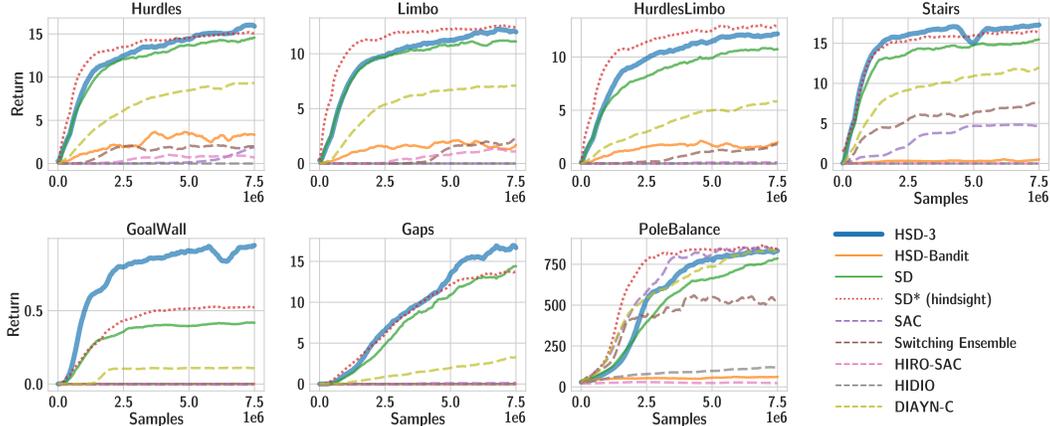


Figure 4: Learning curves on benchmark environments with the Walker robot. For clarity, we plot mean performance over 9 seeds, averaged over 0.5M samples. Full learning curves including error bands are provided in Appendix F.1.

skills. The most general skill, high-level control of *all* considered features (bottom row), also allows for learning progress on all tasks; however, it is generally outperformed by other, more specific skills.

5.2 Skill Selection with HSD-3

We compare HSD-3 against a number of base- and topline methods to examine the effectiveness of our proposed framework for skill learning and hierarchical control. We include state-of-the-art non-hierarchical RL methods (SAC [18]), end-to-end hierarchical RL algorithms (HIRO [33], HIDIO [53]), and evaluate DIAYN as an alternative unsupervised pre-training algorithm [13]. For HIRO, we use SAC for learning both high- and low-level policies (HIRO-SAC). We facilitate the comparison to DIAYN with a continuous-valued skill variable obtained from an embedding [1] (DIAYN-C). For both HIRO and DIAYN-C, we provide the same goal space as in our pre-training stage, consisting of the full feature set, to ensure a fair comparison. We further compare to the Switching Ensemble proposed in [35], which does not learn a high-level policy but has been shown to improve exploration. Finally, we examine how our pre-trained skills perform if we use a single goal space only (SD). Without prior knowledge about the task at hand, this corresponds to the full goal space $F = \{0, 1, \dots, \dim(\mathcal{S}^g)\}$, settling for maximum generality. As a topline, we select the *best per-task* goal spaces from Section 5.1, denoted with SD*, which required exhaustive training and evaluation with all available goal spaces.

In Figure 4, we plot learning curves over 5M training interactions for HSD-3 as well as base- and topline methods, and list final performances in Table 1. On all tasks, HSD-3 outperforms or matches (Limbo) the learning speed of a conventional hierarchical policy with a skill defined over the full goal space (SD). On GoalWall and Stairs, the *three-level policy significantly exceeds the mean performance of single goal space base- and topline*s. For Gaps, the variances of the best results in Table 1 (HSD-3,

Method	Hurdles	Limbo	HurdlesLimbo	Stairs	Gaps	GoalWall	PoleBalance
SAC	2.2±6.2	-0.1±0.2	-0.0±0.4	5.0±4.8	0.1±0.5	-0.2±0.3	866.8±104.5
Switching Ensemble	0.7±3.0	1.9±4.3	1.6±3.6	7.4±3.8	-0.2±0.3	-0.2±0.3	569.4±230.1
HIRO-SAC	0.4±1.6	0.9±2.2	-0.0±0.1	-0.0±0.0	-0.1±0.2	-0.0±0.0	23.8±12.4
HIDIO	-0.1±0.1	-0.1±0.1	-0.2±0.1	-0.2±0.3	-0.2±0.3	-0.3±0.3	117.6±33.8
DIAYN-C	9.5±2.5	7.5±1.3	5.7±1.7	12.5±2.9	3.4±4.1	0.1±0.3	839.9±58.4
SD	15.0±1.4	11.0±1.0	10.8±1.4	15.2±0.9	14.5±9.3	0.4±0.5	789.5±79.1
HSD-Bandit	2.6±1.3	2.4±1.7	2.0±1.5	0.5±0.8	-0.1±0.2	-0.1±0.2	61.0±15.8
HSD-3	15.3±2.0	11.7±0.9	11.8±1.3	17.2±0.7	15.1±8.9	0.9±0.1	876.0±36.9
SD*	15.0±1.0	12.2±1.0	12.7±1.3	16.2±1.3	14.0±11.3	0.5±0.5	868.2±91.1

Table 1: Final performance after 7.5M samples across benchmark tasks with the Walker robot. Mean and standard deviation reported for average returns over 9 seeds.

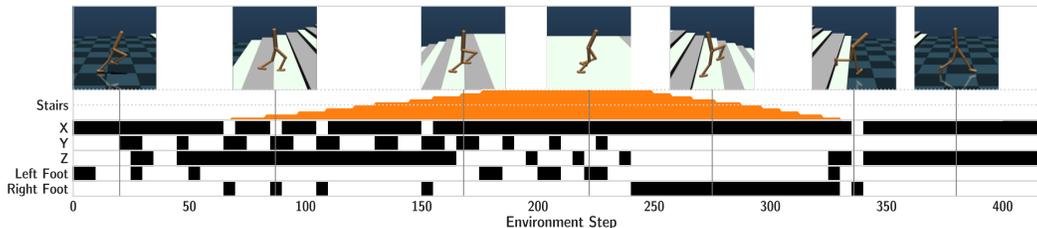


Figure 5: Skills selected by HSD-3 during an episode on the Stairs task. Different skills are utilized for different sections of the environment: walking upstairs is achieved by controlling the torso position (X,Z) in a regular pattern, with goals for the rotation (Y) and the right foot being set occasionally. On the top section, a different pattern is selected for quickly running forwards. When going downstairs, the right foot of the robot is controlled explicitly to maintain balance.

SD, SD*) are high, hinting at especially challenging exploration problems. A closer look at per-seed performance (Appendix F.1) reveals learning failure for 2 out of 9 runs for these methods. For GoalWall, 4 and 5 runs for SD* and SD, respectively, achieve zero return. Per-task goal space selection with a bandit algorithm (HSD-Bandit) fails on most environments, and achieves low returns on Hurdles, Limbo and HurdlesLimbo.

These results demonstrate that, without additional task-specific knowledge, our hierarchical skill framework is able to *automatically and efficiently trade off between skills of varying complexity at training time*. In Figure 5, we visualize the goal spaces that are selected across the course of an episode in the Stairs environment. During different stages of the environment (walking upstairs, downstairs, and on a planar surface), HSD-3 selects appropriate skills of varying complexity².

The performance of the remaining baselines underscores the challenges posed by our benchmark tasks. Standard SAC works well in PoleBalance, which has a dense reward structure compared to the remaining environments. On most other tasks, learning progress is slow or non-existent, e.g., on Hurdles, only 2 out of 9 runs achieve a positive return at the end of training. DIAYN-C exhibits subpar performance across compared our hierarchical pre-training scheme, with the exception of PoleBalance where learning is fast and returns are higher than for SD. The end-to-end HRL methods HIRO-SAC and HIDIO are unable to make meaningful progress on any task, which highlights the utility of skill pre-training in the absence of task-specific priors. The Switching Ensemble is the best baseline without pre-training (except for the PoleBalance task), but a clear majority of individual runs do not obtain positive returns (Appendix F.1). On the Stairs task, both SAC and the Switching Ensemble manage to climb the flight of stairs in several runs, but fail to discover the second, downwards flight at the end of the platform.

We refer to the supplementary material for further ablation and analysis with the Walker robot. In Appendix F.2, we demonstrate the efficacy of our proposed hierarchical pre-training method, compared to pre-training a single skill policy on the full goal space exclusively. In Appendix F.3, we analyze the exploration behavior of various methods from Table 1. We find that, in general, hierarchical methods visit more states compared to SAC. HIRO-SAC and DIAYN-C visit a higher number of states in the Hurdles, Limbo, HurdlesLimbo and Stairs environments but fail to leverage this experience for achieving higher returns.

5.3 Evaluation on a Humanoid Robot

To evaluate the scalability of HSD-3 to more complex robots, we perform an additional set of experiments with the challenging 21-joint Humanoid robot from Tassa et al. [48]. For skill pre-training, we select a goal space approximately matching the Walker setup, consisting of the translation along the X-axis of simulation, Z position of the torso, its rotation around the Y axis and three-dimensional feet positions relative to the robot’s hip location.

In Figure 6, we plot average learning progress across 100M environment steps, comparing HSD-3 to SD (the full goal space), SD* (the best goal space), and a conventional, non-hierarchical SAC setup. The results align with the observations for the Walker robot in Section 5.2 in that HSD-3

²Videos are available at <https://facebookresearch.github.io/hsd3>.

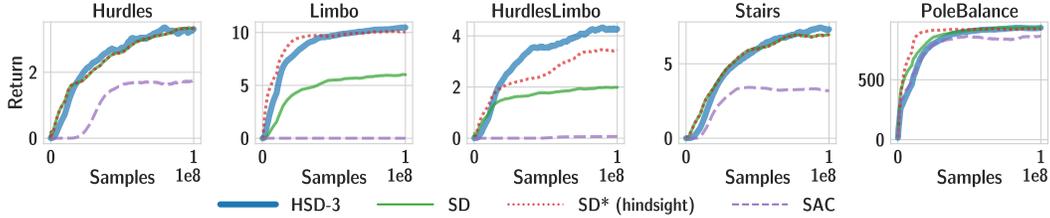


Figure 6: Results on four of the benchmark tasks with a 21-joint Humanoid robot. We show mean performance over 9 seeds, clipped to 0 and averaged over 2.5M samples. Appendix F.5 contains learning curves for all seeds and error bands.

matches or surpasses the performance of the full goal space, and can further outperform the best goal space selected in hindsight. Both HSD-3 and SD make good learning progress on three out of five tasks used for this evaluation (Hurdles, Stairs, PoleBalance); on Hurdles and Stairs, SD is in fact the best single goal space. On HurdlesLimbo, which requires both jumping and crouching, HSD-3 outperforms both SD and SD*. On the Limbo task, HSD-3 is able to learn faster than SD and achieves a higher return by utilizing the most effective single skills on this task, which control X translation and leaning forward or backward. This behavior was not discovered by the SD baseline, which exercises control over all goal space features. In this preliminary study, none of the hierarchical policies was able to make significant progress on the GoalWall and Gaps tasks, which were already shown to be difficult with the Walker robot in Section 5.2.

6 Conclusion

Our work is among the first to highlight the specific trade-offs when embedding prior knowledge in pre-trained skills, and to demonstrate their consequences experimentally in simulated robotic control settings. We describe a hierarchical skill learning framework that, compared to existing approaches, allows high-level policies to perform the trade-off between directed exploration and fine-grained control automatically during training on downstream tasks. Our experiments with a bipedal Walker robot demonstrate HSD-3’s efficacy on a variety of sparse-reward tasks in which previous approaches struggle. We further apply our framework to a challenging Humanoid robot, where it learns effective policies on the majority of tasks.

With this work, we provide a new set of benchmark tasks that require diverse motor skills and pose interesting challenges for exploration. We release the tasks in a dedicated distribution, with the aim of spurring further research on motor skills in continuous control settings, and to ultimately broaden the utility of hierarchical approaches to control to tasks beyond navigation.

Limitations: Our investigations with the Humanoid robot have been performed in limited goal spaces and hence with significant priors. We believe that further work is required to successfully acquire a large, exhaustive set of low-level skills in unsupervised environments that work on more and more complex morphologies.

Acknowledgements: We thank Alessandro Lazaric for insightful discussions, and Franziska Meier, Ludovic Denoyer, and Kevin Lu for helpful feedback on early versions of this paper.

References

- [1] J. Achiam, H. Edwards, D. Amodei, and P. Abbeel. Variational Option Discovery Algorithms. *arXiv:1807.10299 [cs]*, Jul 2018.
- [2] A. Ajay, A. Kumar, P. Agrawal, S. Levine, and O. Nachum. OPAL: Offline Primitive Discovery for Accelerating Offline Reinforcement Learning. In *International Conference on Learning Representations*, 2021.
- [3] P.-L. Bacon, J. Harb, and D. Precup. The Option-Critic Architecture. In *Thirty-First AAAI Conference on Artificial Intelligence*, pages 1726–1734, 2017.
- [4] A. G. Barto, R. S. Sutton, and C. W. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(5):834–846, 1983.
- [5] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. OpenAI Gym. *arXiv:1606.01540 [cs]*, Jun 2016.
- [6] V. Campos, A. Trott, C. Xiong, R. Socher, X. Giro-i-Nieto, and J. Torres. Explore, Discover and Learn: Unsupervised Discovery of State-Covering Skills. In *International Conference on Machine Learning*, pages 1317–1327, 2020.
- [7] M. S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing*, STOC '02, pages 380–388, Montreal, Quebec, Canada, 2002. Association for Computing Machinery.
- [8] P. Dayan and G. E. Hinton. Feudal Reinforcement Learning. In *Advances in Neural Information Processing Systems*, volume 5, pages 271–278. Morgan-Kaufmann, 1992.
- [9] O. Delalleau, M. Peter, E. Alonso, and A. Logut. Discrete and Continuous Action Representation for Practical RL in Video Games. In *AAAI-20 Workshop on Reinforcement Learning in Games*, 2019.
- [10] T. G. Dietterich. Hierarchical Reinforcement Learning with the MAXQ Value Function Decomposition. *Journal of Artificial Intelligence Research*, 13:227–303, 2000.
- [11] P. Dobrowolski. Swing-twist decomposition in Clifford algebra. *arXiv:1506.05481 [cs]*, Jun 2015.
- [12] Y. Duan, X. Chen, R. Houthoofd, J. Schulman, and P. Abbeel. Benchmarking Deep Reinforcement Learning for Continuous Control. In *International Conference on Machine Learning*, pages 1329–1338, 2016.
- [13] 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.
- [14] R. E. Fikes, P. E. Hart, and N. J. Nilsson. Learning and executing generalized robot plans. *Artificial Intelligence*, 3:251–288, 1972.
- [15] C. Florensa, Y. Duan, and P. Abbeel. Stochastic Neural Networks for Hierarchical Reinforcement Learning. In *International Conference on Learning Representations*, 2017.
- [16] D. Ghosh, A. Singh, A. Rajeswaran, V. Kumar, and S. Levine. Divide-and-Conquer Reinforcement Learning. In *International Conference on Learning Representations*, 2018.
- [17] K. Gregor, D. J. Rezende, and D. Wierstra. Variational Intrinsic Control. In *International Conference on Learning Representations (Workshop)*, 2017.
- [18] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, and S. Levine. Soft Actor-Critic Algorithms and Applications. *arXiv:1812.05905v2 [cs]*, Dec 2018.
- [19] K. Hausman, J. T. Springenberg, Z. Wang, N. Heess, and M. Riedmiller. Learning an Embedding Space for Transferable Robot Skills. In *International Conference on Learning Representations*, 2018.

- [20] N. Heess, D. TB, S. Sriram, J. Lemmon, J. Merel, G. Wayne, Y. Tassa, T. Erez, Z. Wang, S. M. A. Eslami, M. Riedmiller, and D. Silver. Emergence of Locomotion Behaviours in Rich Environments. *arXiv:1707.02286 [cs]*, Jul 2017.
- [21] N. K. Jong. The utility of temporal abstraction in reinforcement learning. In *7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 299–306, 2008.
- [22] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015.
- [23] G. Konidaris. Building Portable Options: Skill Transfer in Reinforcement Learning. In *IJCAI*, volume 7, pages 895–900, 2007.
- [24] A. Li, C. Florensa, I. Clavera, and P. Abbeel. Sub-policy Adaptation for Hierarchical Reinforcement Learning. In *International Conference on Learning Representations*, 2019.
- [25] T. Li, R. Calandra, D. Pathak, Y. Tian, F. Meier, and A. Rai. Planning in Learned Latent Action Spaces for Generalizable Legged Locomotion. *arXiv:2008.11867 [undefined]*, 6(2):2682–2689, Mar 2021.
- [26] M. C. Machado, M. G. Bellemare, and M. Bowling. A Laplacian Framework for Option Discovery in Reinforcement Learning. In *International Conference on Machine Learning*, 2017.
- [27] K. Marino, A. Gupta, R. Fergus, and A. Szlam. Hierarchical RL Using an Ensemble of Proprioceptive Periodic Policies. In *International Conference on Learning Representations*, 2019.
- [28] W. Masson, P. Ranchod, and G. Konidaris. Reinforcement Learning with Parameterized Actions. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), 2016.
- [29] J. Merel, Y. Tassa, D. TB, S. Srinivasan, J. Lemmon, Z. Wang, G. Wayne, and N. Heess. Learning human behaviors from motion capture by adversarial imitation. *arXiv:1707.02201 [cs]*, Jul 2017.
- [30] J. Merel, A. Ahuja, V. Pham, S. Tunyasuvunakool, S. Liu, D. Tirumala, N. Heess, and G. Wayne. Hierarchical Visuomotor Control of Humanoids. In *International Conference on Learning Representations*, 2019.
- [31] J. Merel, L. Hasenclever, A. Galashov, A. Ahuja, V. Pham, G. Wayne, Y. W. Teh, and N. Heess. Neural Probabilistic Motor Primitives for Humanoid Control. In *International Conference on Learning Representations*, 2019.
- [32] J. Merel, S. Tunyasuvunakool, A. Ahuja, Y. Tassa, L. Hasenclever, V. Pham, T. Erez, G. Wayne, and N. Heess. Catch & Carry: Reusable Neural Controllers for Vision-Guided Whole-Body Tasks. *ACM Transactions on Graphics*, 39(4), 2020.
- [33] O. Nachum, S. S. Gu, H. Lee, and S. Levine. Data-Efficient Hierarchical Reinforcement Learning. In *Advances in Neural Information Processing Systems*, volume 31, pages 3303–3313. Curran Associates, Inc., 2018.
- [34] O. Nachum, S. Gu, H. Lee, and S. Levine. Near-Optimal Representation Learning for Hierarchical Reinforcement Learning. In *International Conference on Learning Representations*, 2019.
- [35] O. Nachum, H. Tang, X. Lu, S. Gu, H. Lee, and S. Levine. Why Does Hierarchy (Sometimes) Work So Well in Reinforcement Learning? In *NeurIPS 2019 Deep Reinforcement Learning Workshop*, 2019.
- [36] X. B. Peng, M. Chang, G. Zhang, P. Abbeel, and S. Levine. MCP: Learning Composable Hierarchical Control with Multiplicative Compositional Policies. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [37] K. Pertsch, Y. Lee, and J. J. Lim. Accelerating Reinforcement Learning with Learned Skill Priors. In *4th Conference on Robot Learning (CoRL)*, 2020.

- [38] A. H. Qureshi, J. J. Johnson, Y. Qin, T. Henderson, B. Boots, and M. C. Yip. Composing Task-Agnostic Policies with Deep Reinforcement Learning. In *International Conference on Learning Representations*, 2020.
- [39] M. Riedmiller, R. Hafner, T. Lampe, M. Neunert, J. Degraeve, T. Wiele, V. Mnih, N. Heess, and J. T. Springenberg. Learning by Playing - Solving Sparse Reward Tasks from Scratch. In *International Conference on Machine Learning*, pages 4344–4353. PMLR, 2018.
- [40] E. D. Sacerdoti. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence*, 5(2): 115–135, 1974.
- [41] T. Schaul, D. Horgan, K. Gregor, and D. Silver. Universal Value Function Approximators. In *International Conference on Machine Learning*, pages 1312–1320, 2015.
- [42] T. Shankar and A. Gupta. Learning Robot Skills with Temporal Variational Inference. In *International Conference on Machine Learning*, 2020.
- [43] A. Sharma, S. Gu, S. Levine, V. Kumar, and K. Hausman. Dynamics-Aware Unsupervised Discovery of Skills. In *International Conference on Learning Representations*, 2020.
- [44] S. Sinha, H. Bharadhwaj, A. Srinivas, and A. Garg. D2RL: Deep Dense Architectures in Reinforcement Learning. In *NeurIPS 2020 Deep RL Workshop*, 2020.
- [45] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. Adaptive Computation and Machine Learning Series. The MIT Press, Cambridge, Massachusetts, second edition, 2018. ISBN 978-0-262-03924-6.
- [46] 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.
- [47] H. Tang, R. Houthoofd, D. Foote, A. Stooke, X. Chen, Y. Duan, J. Schulman, F. De Turck, and P. Abbeel. #Exploration: A Study of Count-Based Exploration for Deep Reinforcement Learning. In *arXiv:1611.04717 [Cs]*, Advances in Neural Information Processing Systems 30. 2017.
- [48] Y. Tassa, S. Tunyasuvunakool, A. Muldal, Y. Doron, P. Trochim, S. Liu, S. Bohez, J. Merel, T. Erez, T. Lillicrap, and N. Heess. Dm_control: Software and Tasks for Continuous Control. *arXiv:2006.12983 [undefined]*, 6:100022, Nov 2020.
- [49] E. Todorov, T. Erez, and Y. Tassa. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012.
- [50] A. S. Vechnyevets, S. Osindero, T. Schaul, N. Heess, M. Jaderberg, D. Silver, and K. Kavukcuoglu. FeUdal Networks for Hierarchical Reinforcement Learning. In *International Conference on Machine Learning*, pages 3540–3549, 2017.
- [51] E. Wei, D. Wicke, and S. Luke. Hierarchical Approaches for Reinforcement Learning in Parameterized Action Space. In *2018 AAAI Spring Symposia*. AAAI Press, Stanford University, Palo Alto, California, USA, 2018.
- [52] W. Whitney, R. Agarwal, K. Cho, and A. Gupta. Dynamics-aware Embeddings. In *International Conference on Learning Representations*, 2020.
- [53] J. Zhang, H. Yu, and W. Xu. Hierarchical Reinforcement Learning by Discovering Intrinsic Options. In *International Conference on Learning Representations*, 2020.

A Environment Details

For all benchmark tasks, we place the robot at a designated starting position and configuration, and perturb its joint positions with noise sampled uniformly from $[-0.1, 0.1]$ and its joint velocities with noise sampled from $0.1 \cdot \mathcal{N}(0, 1)$. These perturbations are also applied in standard MuJoCo benchmark tasks [5]. Each environment provides three different observations: proprioceptive robot states \mathcal{S}^p , extra task-specific observations \mathcal{S}^+ , and measurements for goal states \mathcal{S}^g . Goal state features are solely used for convenience when providing relative goal inputs to low-level policies (Algorithm 2), and can also be derived from proprioceptive observations and the robot’s absolute position. Below, we list detailed environment configurations and reward functions, as well as robot-specific modifications (if applicable). In all cases, we define invalid states on a per-robot basis (Appendix B) that lead to a premature end of the episode with a reward of -1 . Unless otherwise noted, episodes consist of 1000 steps of agent interaction.

Hurdles Hurdles take the form of simple boxes, placed in intervals $\sim \mathcal{U}(3, 6)$ meters, and with heights $\sim \mathcal{U}(0.1, 0.3)$ meters. Task-specific observations are the distance to the next hurdle and its height. For every hurdle that the robot’s torso passes, the agent receives a reward of 1.

Limbo Cylindrical bars with a diameter of 0.2m are placed in intervals $\sim \mathcal{U}(3, 6)$ meters. Their heights are draw from $\mathcal{U}(1.2, 1.5)$ for the Walker robot, and from $\mathcal{U}(0.9, 1.2)$ for the Humanoid. The agent observes the distance to the next limbo bar, as well as its height above the ground. A reward of 1 is obtained if the agent’s torso position moves past a bar.

HurdlesLimbo This is a combination of the Hurdles and Limbo environments, where hurdles and limbo bars alternate, starting with a hurdle. The reward is defined as above, and the observation consists of the type of obstacle (0 for a hurdle, 1 for a limbo bar) as well as its relative distance and height.

Stairs This task consists of 10 stairs (upwards), a platform, and another 10 stairs (downwards). Stairs have a height of 0.2m, and their length is sampled uniformly within $[0.5\text{m}, 1.0\text{m}]$. The agent observes the distance of the robot’s torso to the next two steps, and a reward of 1 is provided whenever it moves past a step.

Gaps We sample gaps from $\mathcal{U}(0.2, 0.7)$ meters and platforms from $\mathcal{U}(0.5, 2.5)$ meters ($\mathcal{U}(1.0, 2.5)$ meters for the Humanoid). Gaps and platforms are placed in alternation, starting at a distance of 4 meters from the robot’s initial position. Gaps are placed 5cm below the platforms, and the episode ends with a reward of -1 if the robot touches a gap. The agent observes the relative distances to the start of the next gap and next platform. A reward of 1 is achieved if the robot’s torso moves past the start of a platform for the first time.

GoalWall A ball is placed 2.5 meters in front of the robot (1m for the Humanoid). A wall is placed 4 meters from the ball, and contains a circular target area with an 0.8m diameter, placed 1m above ground. For the effectively two-dimensional Walker robot, ball movements are restricted to rotation around the Y-axis and translation along X and Z. At the start of the episode, the ball’s position is perturbed with additive noise drawn from $\mathcal{N}(0, 0.01)$, and we add noise from $\mathcal{N}(0, 0.1)$ to its rotation. The agent observes the position of the ball (X and Y are relative to the robot’s torso, Z is absolute) and its velocity. Episodes end after 250 steps or if the ball’s X position reaches the position of the wall, with a reward of 1 if the ball is inside the target area and 0 otherwise.

PoleBalance A cylindrical pole with a mass of 0.5kg is attached to the robot’s topmost body part. The pole has a length of 0.5m, and the position and velocity of the rotational joint that connects it to the robot is perturbed with additive noise from $\mathcal{N}(0, 0.01)$. For the Walker robot, the pole can rotate around the Y-axis only; for the Humanoid robot, rotations around all axes are permitted. The position and velocity of the joint connecting the pole to the robot are provided as observations. The reward is 1 unless the pole is about to fall over, which terminates the episode. A falling pole is defined by the distance along the Z axis between its lower and upper parts falling below 80% of its length.

In Table 2 we provide overview renderings of the Hurdles, Limbo, HurdlesLimbo, Stairs and Gaps environments to illustrate the positioning of objects along the course. In these tasks, for which locomotion along the X-axis is essential, non-hierarchical baselines agents perform well with a simple shaped reward for movement along the X-axis. For environments with randomly generated obstacles such as ours, the effectiveness of a simple “forward” reward has previously been observed by Heess et al. [20].

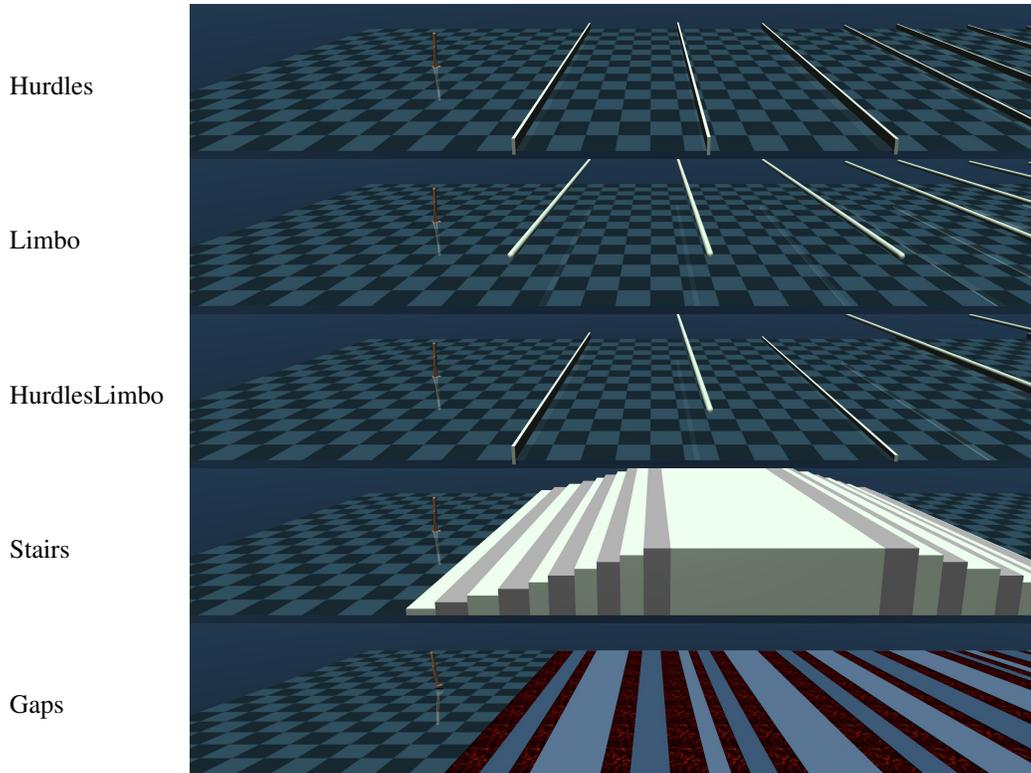


Table 2: Rendering of 5 of our 7 benchmark tasks. Positions of additional objects (hurdles, limbo bars, gaps, stairs) are subject to random perturbation for each episode. Courses for Hurdles, Limbo, HurdlesLimbo and Gaps continue further.

B Robots and Goal Spaces

B.1 Walker

This two-dimensional, bipedal walker has been released as part of the `dm_control` suite [48]. We use an observation space similar to the MuJoCo tasks in Gym [5], where we featurize the positions of all modeled joints, their velocities, and contact forces for all body parts clipped to $[-1, 1]$. We manually construct five primitive goal spaces that can be freely combined to obtain a total of 31 candidate goal spaces:

Feature	Range (min,max)	Direct Obs.
Torso X position	-3, 3	no
Torso Y rotation	-1.3, 1.3	yes
Torso Z position	0.95, 1.5	yes
Left foot, relative to torso	X pos. -0.72, 0.99	no
	Z pos. -1.3, 0	no
Right foot, relative to torso	X pos. -0.72, 0.99	no
	Z pos. -1.3, 0	no



Directly observable features are included in proprioceptive observations; non-observable ones can either be derived (i.e., the foot position can be determined via joint positions) or correspond to global observations (X position). All position features are measured in meters while rotations are expressed in radians. Ranges were obtained via rollouts from standard SAC policies trained with shaping rewards.

For the Walker robot, we define invalid states as the torso’s Z position falling below 0.9 meters, or its Y rotation having a value outside of $[-1.4, 1.4]$. These limits are intended to prevent the robot from falling over.

B.2 Humanoid

We adopt the simplified Humanoid robot from Tassa et al. [48], which consists of 21 joints. We use a modified observation space with the same type of state features as the Walker robot (B.1). Following recent work on skill learning for Humanoid robots [36, 31], actions consist of target joint positions, mapped to $[-1, 1]$. We found this to produce better results compared to specifying raw torques.

For goal space featurization, closely follow our Walker setup and provide torso X and Z positions, Y rotation and relative foot positions. Feature values for rotations around Y and Z were obtained with a twist-and-swing decomposition of the torso’s global orientation matrix [11]. State features for learning agents use a standard quaternion representation, however.

Feature	Range (min,max)	Direct Obs.
Torso X position	-3, 3	no
Torso Y rotation	-1.57, 1.57	no
Torso Z position	0.95, 1.5	yes
Torso Z rotation	-1.57, 1.57	no
Left foot, relative to hip	X pos.	-1, 1
	Y pos.	-1, 1
	Z pos.	-1, 0.2
Right foot, relative to hip	X pos.	-1, 1
	Y pos.	-1, 1
	Z pos.	-1, 0.2



By coupling every feature combination with the Z rotation feature, we again obtain $2^5 - 1 = 31$ possible goal spaces as for the Walker robot.

Similar to the Walker robot, we define valid states as the torso’s position being at least 0.9 meters above ground. Body rotations are not restricted.

B.3 Goal Space Construction

Our main considerations for formalizing goal spaces are the elimination of bias due to different feature magnitudes and a convenient notation for considering subsets of features from the goal feature space \mathcal{S}^g . Starting from a set of $n := \dim(\mathcal{S}^g)$ features with ranges $(l^i, h^i) : 1 \leq i \leq n$, we construct a goal space transformation matrix Ψ along with offsets b :

$$\Psi := 2I_n \begin{bmatrix} (h^1 - l^1)^{-1} \\ (h^2 - l^2)^{-1} \\ \vdots \\ (h^n - l^n)^{-1} \end{bmatrix}, \quad b := -2 \begin{bmatrix} l^1 (h^1 - l^1)^{-1} \\ l^2 (h^2 - l^2)^{-1} \\ \vdots \\ l^n (h^n - l^n)^{-1} \end{bmatrix} - 1$$

A single goal space over features $F = \{i, j, k, \dots\}$ is defined as the image of an affine function $\omega^F : \mathcal{S}^g \rightarrow \mathcal{G}^F$,

$$\omega^F(s) := s \begin{bmatrix} \Psi_{i,*} \\ \Psi_{j,*} \\ \Psi_{k,*} \\ \vdots \end{bmatrix}^T + \begin{bmatrix} b_i \\ b_j \\ b_k \\ \vdots \end{bmatrix}$$

ω^F can be understood as simple abstraction that selects a subset F of goal space features and applies a suitable normalization, mapping $[l^i, h^i]$ to $[-1, 1]$ for $i \in F$. A desirable effect of this normalization is that distance-based rewards that are commonly used to train goal-based policies will also be normalized, which facilitates optimization when sharing the parameters of policies across several goal spaces. Additionally, the resulting high-level action space for subgoals can be conveniently defined over $[-1, 1]^F$.

C Unsupervised Pre-Training Details

Our unsupervised pre-training algorithm is provided in Algorithm 1. We assume that the pre-training environment provides access to both proprioceptive states (the input of the skill policy) and goal state features as defined in Appendix B. During training, goal spaces and goals are randomly selected for each episode. The low-level skill policy π^{lo} and the corresponding Q-functions are trained with the standard Soft Actor-Critic update [18, Alg. 1], with representations of both goal features F and goal g considered a part of the state observations.

The STEP_ENV function in Algorithm 1 samples a transition from the pre-training environment, computes the reward as described in Section 3.2 and determines the values of done and reset. done is set to true if one of the following conditions hold:

- The current goal is reached (i.e., the L2 distance to it is smaller than ϵ_g).
- A fixed number of steps (horizon h) has passed since the previous goal was sampled.
- A random number sampled from $[0, 1)$ is smaller than the resample probability ϵ_r .

reset is set to true if any of these conditions hold:

- The state that was reached is considered invalid (A). In this case, the reward is set to -1 .
- A fixed number of goals (reset interval n_r) have been sampled without a simulation reset.

Algorithm 1 Unsupervised Pre-Training of Hierarchical Skills, see Figure 1 for context

Require: Goal spaces defined via feature sets \mathcal{F} and transformations ω^F

```

1: Pre-training environment  $E$  with  $\mathcal{S} = \mathcal{S}^p \cup \mathcal{S}^g$ 
2: Initialize policy  $\mu_\theta : \mathcal{S}^p \times \mathcal{F} \times \mathcal{G} \rightarrow \mathcal{A}$ 
3: Initialize Q-function  $Q_{\phi_i} : \mathcal{S}^p \times \mathcal{F} \times \mathcal{G} \times \mathcal{A} \rightarrow \mathbb{R}$  for  $i \in \{1, 2\}$ 
4: Replay buffer  $B \leftarrow \text{CIRCULARBUFFER}()$ 
5: reset  $\leftarrow$  true
6: for  $i \leftarrow 1, N$  do
7:   if reset then
8:      $s \sim S_0$  ▷ Reset simulation
9:   end if
10:   $F \sim \mathcal{F}, g \sim [-1, 1]^F$  ▷ Sample new goal space and goal
11:  reset  $\leftarrow$  false, done  $\leftarrow$  false
12:  while not (done or reset) do
13:     $a \sim \mu_\theta(s^p, F, (\omega^F)^{-1}(g) - s^g)$ 
14:     $s', r, \text{done}, \text{reset} \leftarrow \text{STEP\_ENV}(s, a, F, g)$ 
15:     $B.append(s^p, F, (\omega^F)^{-1}(g) - s^g, a, r, s')$ 
16:     $s \leftarrow s'$ 
17:    if  $i \% f_u == 0$  then
18:      for each gradient step do
19:         $\phi, \theta \leftarrow \text{SAC\_UPDATE}(\phi, \theta, B)$  ▷ Perform standard SAC update
20:      end for
21:    end if
22:  end while
23: end for
24: Output:  $\theta, \phi_1, \phi_2$ 

```

D Hierarchical Control Details

D.1 Soft-Actor Critic for HSD-3

Below, we provide explicit derivations for extending Soft Actor-Critic to a factorized action space $\mathcal{F} \times \mathcal{G}$ for π^{hi} , consisting of discrete actions $F \in \mathcal{F}$ (goal space feature sets) and continuous actions $g \in \mathcal{G}^F$ (goals). Our extension is performed according to the following desiderata: (1) We utilize a shared critic $Q(s, F, g)$ and two policies $\pi^f : \mathcal{S} \rightarrow \mathcal{F}, \pi^g : \mathcal{S} \times \mathcal{F} \rightarrow \mathcal{G}^F$; (2) we compute π^f for all

Parameter	Value (Walker)	Value (Humanoid)
Optimizer	Adam [22]	Adam
Learning rate λ_Q	0.001	0.001
Learning rate λ_π	0.001	0.001
Learning rate λ_α	0.001	0.001
Target entropy $\overline{\mathcal{H}}$	$-\dim(\mathcal{A}) = -6$	$-\dim(\mathcal{A}) = -21$
Initial temperature α	0.1	0.1
Target smoothing coefficient τ	0.005	0.005
Control cost ζ	0.01	0
Horizon h	72	72
Discount factor γ	$1 - 1/h$	$1 - 1/h$
Goal threshold ϵ_g	0.1	0.1
Resample probability ϵ_r	0	0.01
Reset interval n_r	100	100
Replay buffer size	$3 \cdot 10^6$	$3 \cdot 10^6$
Parallel environments	20	40
Environment steps per iteration	1000	5000
Gradient steps per iteration	50	50
Mini-batch size	256	1024
Warmup steps	10^4	10^4
Total iterations	10^4	$3.7 \cdot 10^4$

Table 3: Hyper-parameters for unsupervised pre-training.

discrete actions, which provides us with better value estimates a richer policy learning signal; (3) separate, automatically adjusted temperature coefficients are used for the discrete-action policy (α) and continuous-action policy (β). As the action space of π^g is conditioned on F , we use separate coefficients β^F ; (4) we further normalize the entropy contributions from the different action spaces \mathcal{G}^F by their dimensionality $|F|$, computing action log-probabilities as $|F|^{-1} \log \pi^g(g|s, F)$.

The soft value function $V(s)$ [18, Eq. 3] for acting with both policies π^f and π^g is given as follows:

$$V(s) = \mathbb{E}_{\substack{F \sim \pi^f(\cdot|s), \\ g \sim \pi^g(\cdot|s, F)}} \left[Q(s, F, g) - \alpha \log \pi^f(F|s) - \frac{\beta^F}{|F|} \log \pi^g(g|s, F) \right]$$

Computing the expectation over discrete actions F explicitly yields

$$\begin{aligned} V(s) &= \sum_{F \in \mathcal{F}} \pi^f(F|s) \left(\mathbb{E}_{g \sim \pi^g} \left[Q(s, F, g) - \frac{\beta^F}{|F|} \log \pi^g(g|s, F) \right] - \alpha \log \pi^f(F|s) \right) \\ &= \sum_{F \in \mathcal{F}} \pi^f(F|s) \mathbb{E}_{g \sim \pi^g} \left[Q(s, F, g) - \frac{\beta^F}{|F|} \log \pi^g(g|s, F) \right] + \alpha \mathcal{H}(\pi^f(\cdot|s)) \end{aligned}$$

We arrive at the formulation in Section 3.3 by subtracting the entropy of the uniform discrete-action policy, $\log |\mathcal{F}|$, from $\mathcal{H}(\pi^f(\cdot|s))$ to ensure negative signage. We proceed in an analogous fashion for policy [18, Eq. 7] and temperature [18, Eq. 18] losses.

The resulting high-level policy training algorithm is listed in Algorithm 2. For brevity, loss and value function definitions above use a single variable for states which indicates both proprioceptive observations s^p and task-specific features s^+ .

E Training Details

Neural Network Architecture For all experiments, we use neural networks with 4 hidden layers, skip connections and ReLU activations [44]. Neural networks that operate on multiple inputs (such as the skill policy, or the Q-function in HSD-3) are provided with a concatenation of all inputs. For the Q-function in HSD-3 and HSD-Bandit, goal inputs occupy separate channels for each goal space. The goal policy in HSD-3 π^g is modelled as a multi-head network to speed up loss computations.

Algorithm 2 HSD-3 High-level Policy Training, see Figure 1 for context

Require: Low-level policy π_θ^{lo}
Require: Goal spaces defined via feature sets \mathcal{F} and transformations ω^F
Require: High-level action interval $c \in \mathbb{N}^+$

- 1: Initialize policies π_ϕ^f, π_ψ^g , Q-functions Q_{ρ_i} $i \in \{1, 2\}$, target networks $\bar{\rho}_1 \leftarrow \rho_1, \bar{\rho}_2 \leftarrow \rho_2$
- 2: Replay buffer $B \leftarrow \text{CIRCULARBUFFER}()$
- 3: $s \sim S_0, t \leftarrow 0$
- 4: **for** each iteration **do**
- 5: **for** each environment step **do**
- 6: **if** $t \% c == 0$ **then** ▷ Take high-level action
- 7: $F \sim \pi_\phi^f([s^p; s^+])$
- 8: $g \sim \pi_\psi^g([s^p; s^+], F)$
- 9: $g' \leftarrow (\omega^F)^{-1}(g) - s^g$ ▷ Backproject goal to \mathcal{S}^g and obtain delta
- 10: **end if**
- 11: $a \leftarrow \mathbb{E}_a [\pi_\theta^{lo}(a|s^p, F, g')]$ ▷ Act with deterministic low-level policy
- 12: $s' \sim p_E(s'|s, a)$ ▷ Sample transition from environment
- 13: $B.append(s, F, g, a, t, r(s, a), s')$
- 14: $g' \leftarrow s^g - s'^g + g'$ ▷ Update goal
- 15: $s \leftarrow s'$
- 16: $t \leftarrow t + 1$
- 17: **if** end of episode **then**
- 18: $s \sim S_0, t \leftarrow 0$
- 19: **end if**
- 20: **end for**
- 21: **for** each gradient step **do**
- 22: Sample mini-batch from B , compute losses ▷ See Section 3.3
- 23: Update network parameters $\phi, \psi, \rho_1, \rho_2$
- 24: Update temperatures α, β
- 25: Update target network weights $\bar{\rho}_1, \bar{\rho}_2$
- 26: **end for**
- 27: **end for**
- 28: **Output:** $\psi, \phi, \rho_1, \rho_2$

Rather than receiving an input for the selected current goal space, the output of the respective head is selected.

Hyper-Parameters are listed in Table 3 for pre-training and Table 5 for HSD-3 high-level policy training. Downstream task training runs with fixed, single goal spaces (SD, SD*) use identical hyper-parameters, but do not require $\lambda_f, \bar{\mathcal{H}}^f$, and α . Baseline runs use a learning rate of 0.003 for neural networks and an initial temperature of 0.1 [44]. For HSD-3 and SD, we searched for learning rates in $\{0.0001, 0.0003\}$ and initial temperature values in $\{0.1, 1\}$.

Evaluation Protocol In regular intervals (Table 5), we perform 50 trials with a deterministic policy and measure the average return that was achieved across trials. Since initial states and environment configurations are randomly sampled, we use the same set of 50 environment seeds for all evaluations to ensure comparability.

E.1 Baselines

E.1.1 HIRO-SAC

We implement an end-to-end HRL method similar to HIRO [33], but with Soft Actor-Critic as the underlying learning algorithm rather than TD3. High-level policy actions are expressed within the goal space that we use for pre-training skill policies (B). The neural network policies we use for SAC predict both mean and variances of a Gaussian distribution, and hence we perform goal relabelling by maximizing the log-probabilities of low-level actions directly instead of approximating them via squared action differences [33, Eq. 5]. We observed worse performance by including DynE critic

updates [52]. We therefore report results with an update schedule similar to [33], where high-level policies are being updated less frequently than low-level ones (depending on the high-level action frequency).

E.1.2 DIAYN-C

In a setup similar to Achiam et al. [1], we learn a continuous representation of DIAYN’s discrete skill variable. We input the one-hot encoded skill index to a linear layer with 7 outputs, which corresponds to $\dim(S^g)$ for the walker. Its output is subject to a hyperbolic tangent activation so that the final low-level policy skill input is bounded in $[-1, 1]$. We operate DIAYN’s discriminator on our pre-defined goal spaces (B) and hence provide the same prior knowledge as in our methods. After pre-training with DIAYN, we train a high-level policy as we do for SD baselines, providing its actions directly to the pre-trained policy. We ran experiments with 256 or 1024 hidden units for the skill policy, and with 5,10,20,50 and 100 discrete skills for pre-training. We found that best overall performance was achieved with 256 hidden units for the low-level policy and 10 different skills.

E.1.3 Switching Ensemble

As proposed by Nachum et al. [35], this baseline consists of a small number of standard policies that gather shared experience and are randomly selected for short time-spans during rollouts. In our version, we use SAC as the underlying learning algorithm, in the same configuration as for the SAC baseline. We use an ensemble of 5 policies, and perform switches with the same frequency that high-level actions are taken at for the other hierarchical methods. For evaluations, we act with a single policy throughout the entire episode.

E.1.4 HIDIO

We use the official implementation from Github³. We found it non-trivial to present the low-level policy discriminator with observations in our standard goal space. Hence, in contrast to HIDIO-SAC and DIAYN-C, the HIDIO baseline discriminator operates on the full state space. In accordance with the other hierarchical methods considered, the steps per option were set to 1 on PoleBalance and 5 otherwise, and the discount factor was set to 0.99. Likewise, we use 256 hidden units for the low-level policy’s neural network layers. We performed a hyper-parameter sweep on the Hurdles task, similar to the one performed in the original paper [53] (Table 4).

Parameter	Value
Discriminator input	state, action, state_action, state_difference
Latent option vector dimension (D)	8 , 12
Rollout length	25, 50, 100
Replay buffer length	50000, 200000

Table 4: Hyper-parameters considered for HIDIO, with best ones emphasized.

³<https://github.com/jesbu1/hidio/tree/245d758>

Parameter	Value (Walker)	Value (Humanoid)
Optimizer	Adam [22]	Adam
Learning rate λ_Q	0.001	0.001
Learning rate λ_f	0.003	0.001
Learning rate λ_g	0.003	0.003
Learning rate λ_α	0.001	0.001
Learning rate λ_β	0.001	0.001
Target entropy $\overline{\mathcal{H}}^f$	$0.5 \log \mathcal{F} $	$0.5 \log \mathcal{F} $
Target entropy $\overline{\mathcal{H}}^g$	-1	-1
Initial temperatures α, β^F	1	1
Target smoothing coefficient τ	0.005	0.005
Discount factor γ	0.99	0.99
Replay buffer size	10^6	$2 \cdot 10^6$
Parallel environments	1	5
Environment steps per iteration	50	500
Gradient steps per iteration	50	50
Mini-batch size	256	512
Warmup steps	10^3	10^4
Evaluation interval (iterations)	1000	400

Table 5: Hyper-parameters for high-level policy training with HSD-3.

F Extended Results

Below, we provide full learning curves for the results presented in Section 5 and additional experiments.

F.1 Walker Learning Curves

Learning curves for baselines, HSD-3, HSD-Bandit and SD are provided in Table 6. In addition to the discussion in the experimental section, these plots emphasize that exploration is challenging in all environments apart from PoleBalance. For SAC, SE and HIRO-SAC in particular, positive returns are obtained with few seeds only (gray lines).

F.2 Ablation: Multi-Task Pre-Training

We perform an ablation study on the impact of our multi-task pre-training algorithm used to obtain a hierarchy of skills (C). We train SD high-level policies, i.e., with a goal space consisting of all considered features, with a skill policy that was trained to reach goals defined in this goal space only. This is in contrast to the pre-trained models used in the other experiments throughout the paper, which are trained to achieve goals within a hierarchy of goal spaces. Networks used to train the single goal-space skill policy consist of 256 units per hidden layer, while skill policies shared among multiple goal spaces use 1024 units (increasing the number of parameters for the single skill policy resulted in worse performance). The results in Figure 7 show that shared policies obtained with multi-task pre-training yield higher returns in most benchmark environments. For GoalWall, the usage of a single-skill policy prevented learning progress altogether. These findings indicate the effectiveness of the multi-task pre-training stage: it not only produces a compact representation of various skills, making downstream usage more practical, but also results in improved policies for individual goal spaces.

F.3 Analysis: Exploration Behavior

For gaining additional insight into how HSD-3 impacts exploration behavior, we analyze state visitation counts over the course of training. Due to the continuous state spaces of our environments, we estimate the number of unique states visited over the course of training with SimHash, a hashing method originally proposed for computing exploration bonuses [7, 47]. We hash the full observation, which can include other, randomly placed objects (such as Hurdles or a ball). We compare the amount of unique hashed states for HSD-3 and selected baselines with the Walker robot in Figure 8.

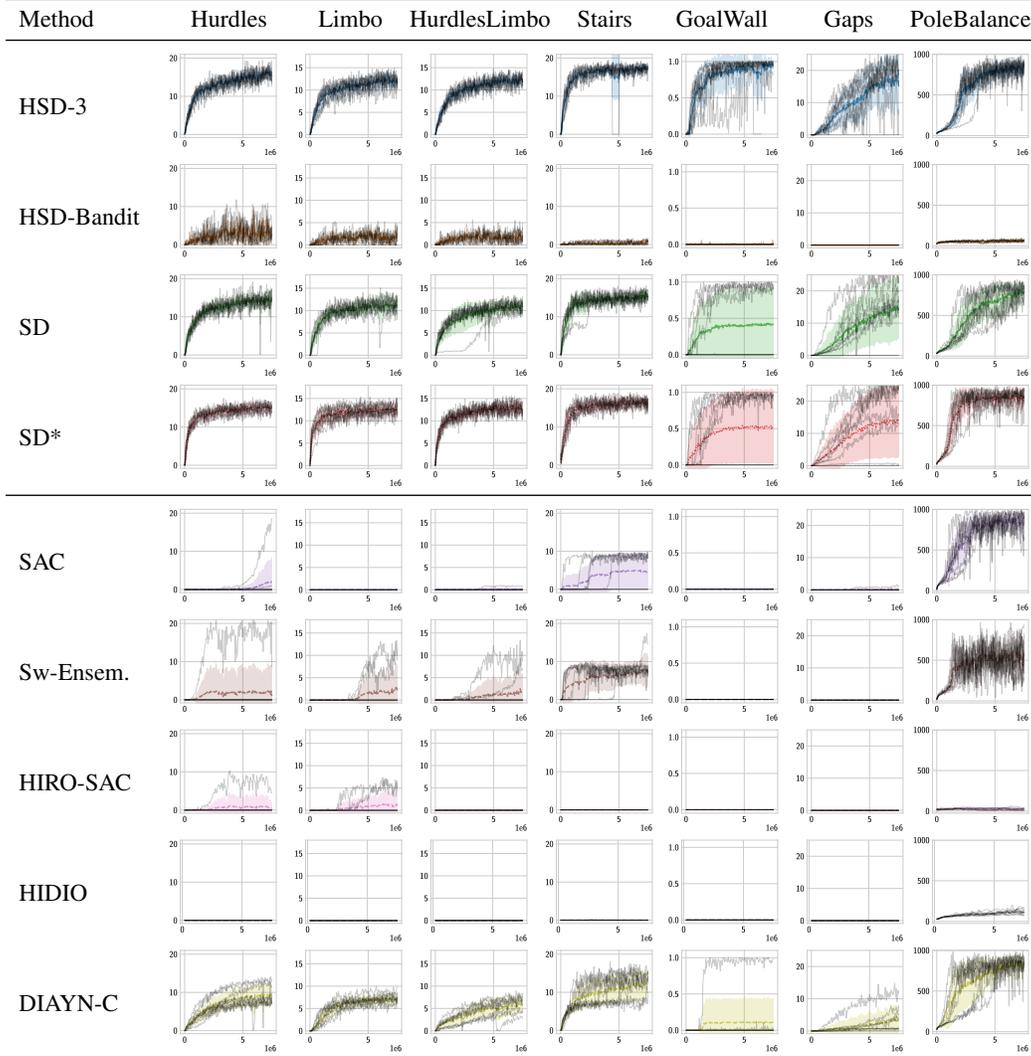


Table 6: Full learning curves for results reported in Table 1 (Walker). We show mean returns achieved (Y axis) after interacting with an environment for a given number of steps (X axis). Shaded areas mark standard deviations computed over 9 seeds for each run.

Generally, hierarchical methods encounter more states during training compared to SAC, even if this does not necessarily translate to a higher return (cf. Figure 4). For HSD-3, after an initial phase of fast learning in Hurdles, Limbo, HurdlesLimbo and Stairs, the performance in terms of return converges, which is reflected in a decrease of new states being visited.

F.4 Analysis: Pre-Training Performance

In Figure 9, we plot performance during unsupervised pre-training. We train a set goal-reaching policies for different feature combinations, modeled with a single neural network. We found that in our pre-training setup, the training reward alone does not adequately capture the fraction of goals that can be reached reliably. As the number of feature subsets increases, dedicated evaluation runs needlessly prolong the wall-clock training time. As an alternative measure of performance, we train an additional Q-function on a 0/1 reward (1 if a goal was reached) and query it periodically with initial states from the replay buffer and randomly sampled goals. The resulting metric reflects controllability, i.e., the probability with which a randomly sampled goal can be reached.

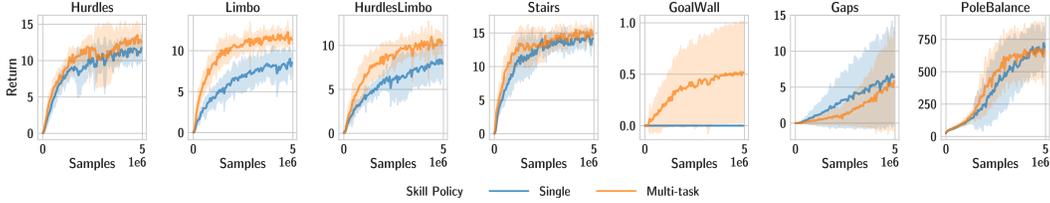


Figure 7: Returns achieved after 5M samples with the SD baseline and the Walker robot, using low-level policies that were trained on a single goal space or with our proposed multi-task pre-training scheme.

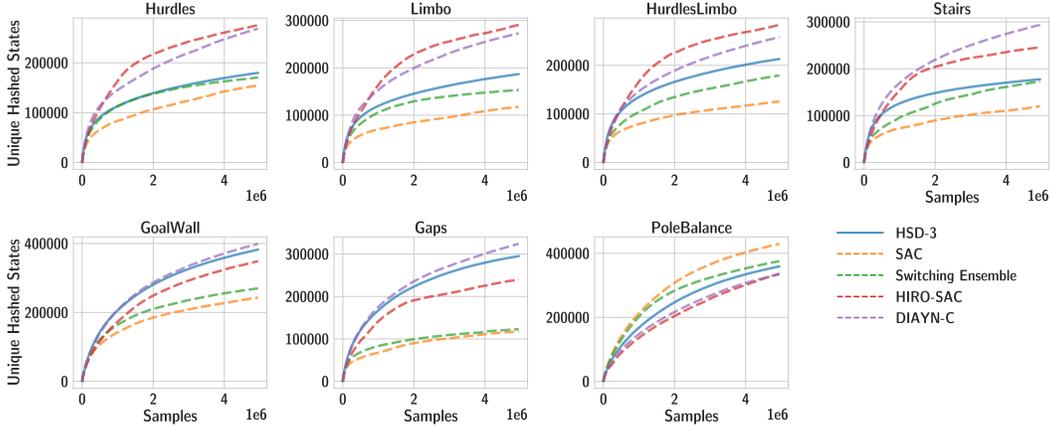


Figure 8: Number of unique hashed states encountered over the course of training (Walker). Mean over 9 seeds per task and method.

Figure 9 shows that, with an increasing number of features, goals become harder to reach. While it is unsurprisingly harder to achieve goals in many dimensions, another effect is that we do not account for fundamentally unreachable states in our definition of the goal space S^g . For example, the reachable states for the feet (LF, RF), which are two-dimensional features (X and Z position) roughly describe a half moon while we sample goals from a hypercube. This effect is multiplied when feature sets are combined.

F.5 Humanoid Learning Curves

In Figure 10, we plot performance for high-level policy training with fixed, individual goal spaces on the Humanoid robot. Similar to the results for the Walker robot (Figure 3), the best goal space features differ significantly across tasks.

Learning curves with the Humanoid robot are presented in Table 7. The non-hierarchical SAC baseline achieves positive returns in the Stairs environment only. The two successful runs manage to climb up the initial flight of stairs but fall over at the top, thus not making it through the entire course. The mediocre median performance of HSD-3 on Stairs is attributed to one of the three seeds failing to learn; the other two are on par with the SD baseline. Notably, HSD-3 outperforms the best single goal space on the Limbo and HurdlesLimbo tasks.

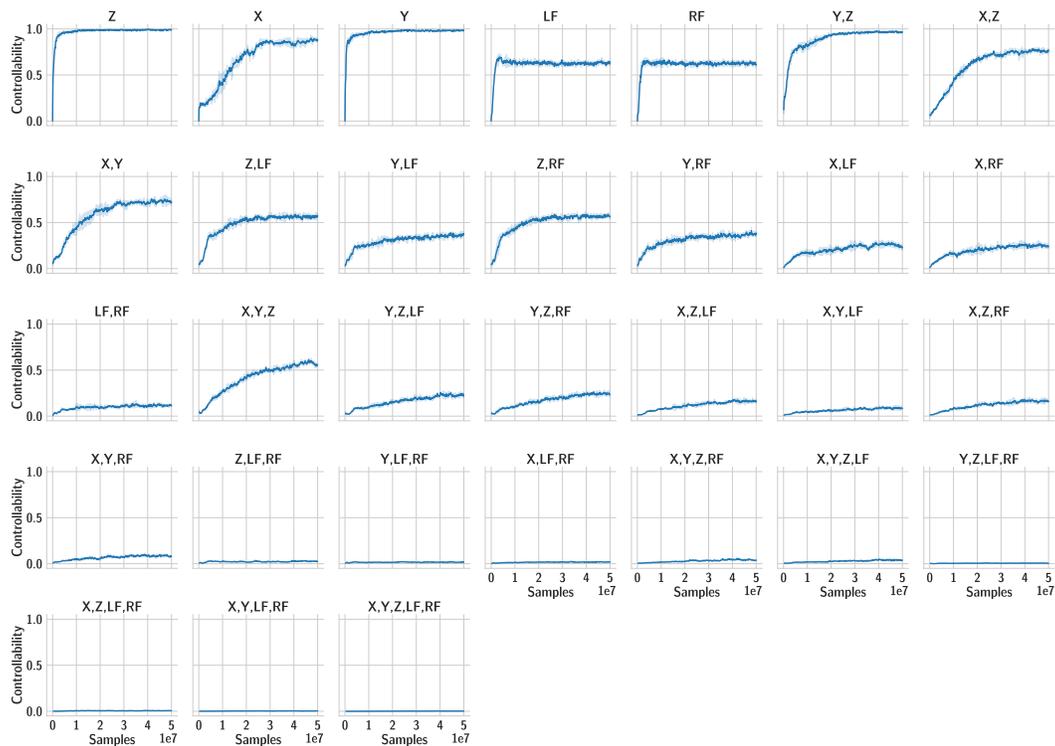


Figure 9: Pre-training performance over the different feature sets considered (Walker robot). Controllability (Y axis) is estimated with a dedicated model. Mean and standard deviation over 3 runs.

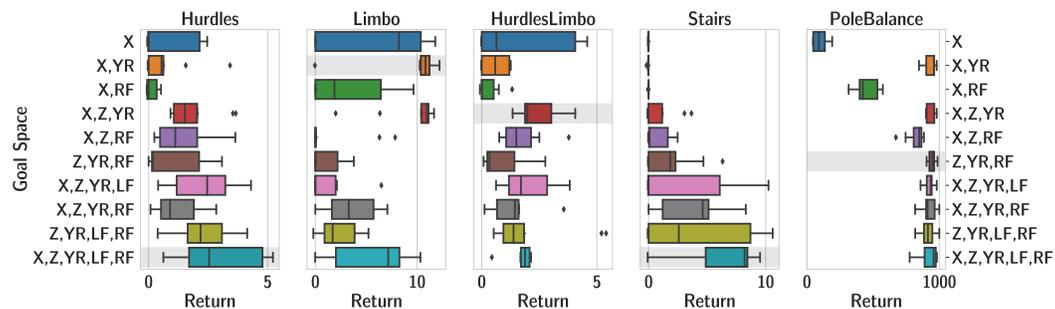


Figure 10: Returns achieved after 50M samples on the benchmark tasks with the Humanoid robot with fixed low-level policy goal spaces (individual runs and quartiles). Each row corresponds to a set of features for the respective goal space. All goal spaces also include the Z rotation of the torso as a feature.

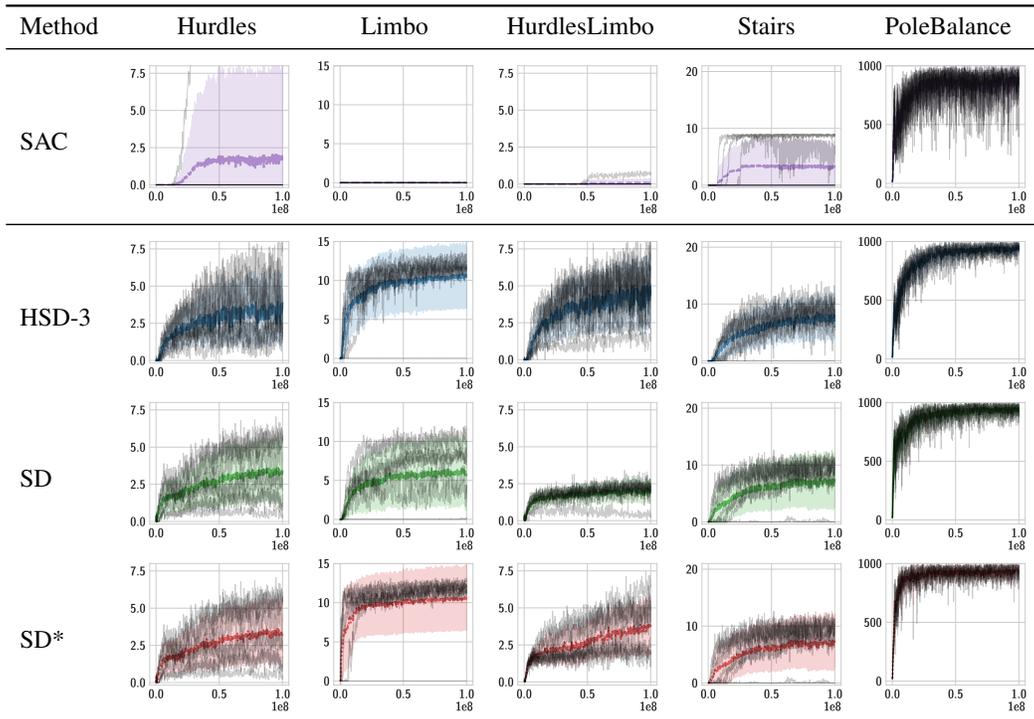


Table 7: Full learning curves for results reported in Figure 6 (Humanoid). We show mean returns achieved (Y axis) after interacting with an environment for a given number of steps (X axis). Shaded areas mark standard deviations computed over 9 seeds for each run.

G Implicit Priors in HIRO

To highlight the prevalence of implicit priors towards navigation tasks in existing methods, we perform an ablation study on HIRO, a prominent end-to-end HRL method [33, 34]. HIRO defines manual subgoal ranges to steer its low-level policy, where ranges for X and Y locations are of higher magnitude compared to joint angles. In Figure 11, we analyze the impact of normalizing the reward on the AntMaze task. We compute the intrinsic reward for the low-level policy [33, Eq.3] as

$$r(s_t, g_t, a_t, s_{t+1}) = - \left\| \frac{s_t - g_t + s_{t+1}}{R} \right\|_2,$$

with $R = [10, 10, 0.5, 1, 1, 1, 1, 0.5, 0.3, 0.5, 0.3, 0.5, 0.3, 0.5, 0.3]$ corresponding to the high-level policy’s action scaling [33, C.1]. The normalization causes all dimensions of the goal space to contribute equally to the low-level policy’s learning signal, and effectively removes any implicit prior towards center-of-mass translation. As a result, among the three goals used for evaluation, only the nearby one at position (16,0) can be successfully reached without reward normalization.

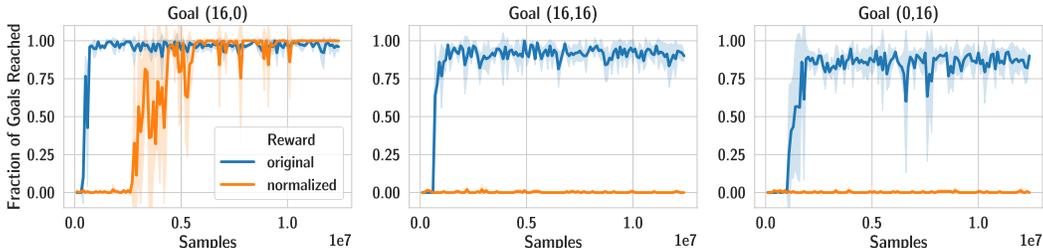


Figure 11: Fraction of goals reached in the AntMaze task with HIRO during evaluations (mean and standard deviation over 3 seeds). With normalized rewards, locomotion on the X-Y plane is no longer subject to higher rewards, and only the nearby goal at position (16,0) can be reached.

Nachum et al. [34] propose to automatically learn a state representation used as the goal space for high-level actions. We investigate the dependence of inductive bias towards X-Y translation by modifying their setup as follows. First, we limit the input to the state representation function f [34, Sec. 2] to the same 15 features that serve as a goal space in [33]. We then normalize the input of f as above via division by R , and modify the high-level action scaling accordingly (actions in $[-1, 1]^2$, and a Gaussian with standard derivation 0.5 for exploration). As a result, the agent is unable to reach any of the tested goals (Figure 12). The limitation of input features alone does not affect performance (blue curve).

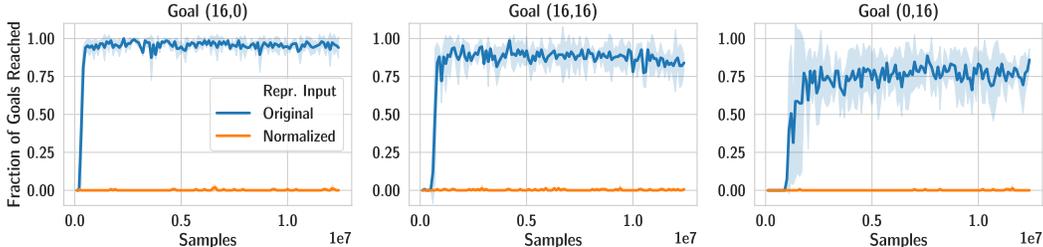


Figure 12: Fraction of goals reached in the AntMaze task with HIRO and representation learning during evaluations (mean and standard deviation over 3 seeds). With normalized state representation inputs, no learning progress is made.