

SkillTree: Explainable Skill-Based Deep Reinforcement Learning for Long-Horizon Control Tasks

Yongyan Wen¹, Siyuan Li^{1*}, Rongchang Zuo¹, Lei Yuan^{2,3,4}, Hangyu Mao⁵, Peng Liu¹

¹Faculty of Computing, Harbin Institute of Technology

²National Key Laboratory of Novel Software Technology, Nanjing University

³School of Artificial Intelligence, Nanjing University

⁴Polixir Technologies

⁵Kuaishou Technology

23B903027@stu.hit.edu.cn, siyuanli@hit.edu.cn

Abstract

Deep reinforcement learning (DRL) has achieved remarkable success in various domains, yet its reliance on neural networks results in a lack of transparency, which limits its practical applications in safety-critical and human-agent interaction domains. Decision trees, known for their notable explainability, have emerged as a promising alternative to neural networks. However, decision trees often struggle in long-horizon continuous control tasks with high-dimensional observation space due to their limited expressiveness. To address this challenge, we propose SkillTree, a novel hierarchical framework that reduces the complex continuous action space of challenging control tasks into discrete skill space. By integrating the differentiable decision tree within the high-level policy, SkillTree generates discrete skill embeddings that guide low-level policy execution. Furthermore, through distillation, we obtain a simplified decision tree model that improves performance while further reducing complexity. Experiment results validate SkillTree’s effectiveness across various robotic manipulation tasks, providing clear skill-level insights into the decision-making process. The proposed approach not only achieves performance comparable to neural network based methods in complex long-horizon control tasks but also significantly enhances the transparency and explainability of the decision-making process.

Code — <https://github.com/YongyanWen/SkillTree>

Introduction

Deep Reinforcement learning (DRL) has been shown as a powerful framework for tackling complex decision-making tasks, achieving remarkable success in various domains such as games (Mnih et al. 2015; BAAI 2023), robotic manipulation (Akkaya et al. 2019; Jitoshio et al. 2023), and visual navigation (Kulhánek, Derner, and Babuška 2021). Despite these advancements, the black-box nature of neural networks poses significant challenges in understanding and trusting their decision making processes. This lack of transparency is particularly concerning in safety-sensitive and human-agent interaction applications, where understanding

the rationale behind decisions is essential (Hickling et al. 2023). For example, in the autonomous driving domain, only if human users can understand the driving policies can they trustfully use them in their cars.

Explainable reinforcement learning (XRL) (Heuillet, Couthouis, and Díaz-Rodríguez 2021; Hickling et al. 2023; Milani et al. 2024) aims to enhance the transparency and explainability of DRL models. XRL methods can be broadly categorized into inherently explainable models, where explainability is induced in the model training process, and post-hoc explanations, where models are explained after training. Existing works have shown that using a decision tree (DT) (Costa and Pedreira 2023) as the underlying model is effective (Frosst and Hinton 2017; Bastani, Pu, and Solar-Lezama 2018; Ding et al. 2020; Vasić et al. 2022). In a DT, observations are input at the root node, and decisions are determined by conditional branches leading to the leaf nodes, which provide the final outputs. This simple structure gives DTs high explainability, offering clear and straightforward decision paths. Unlike post-hoc methods that rely on external algorithms to interpret black-box models, DT-based methods embed transparency directly within the model. The DT-based XRL paradigm not only enhances clarity but also simplifies debugging and validation, as the rationale behind each decision is explicitly encoded within the tree.

Despite their advantages, DT-based methods are not suitable for the following challenging task settings. (a) **Long-horizon tasks**: The temporally-extended decision processes require large and complex trees, which are difficult to optimize (Liu et al. 2021). (b) **High-dimensional state spaces**: DTs often lack sufficient representational ability to effectively manage high-dimensional state spaces, leading to sub-optimal performance in these environments (Bastani, Pu, and Solar-Lezama 2018; Ding et al. 2020). (c) **Continuous action spaces**: The limited number of leaf nodes constrains DTs’ ability to encode continuous control policies optimally, particularly in complex robotic tasks. These limitations restrict the applicability of DT-based RL methods in complex environments, highlighting the need for approaches that can manage high-dimensional spaces and long-horizon tasks.

To address the limitations of traditional DTs in handling long-horizon, high-dimensional, and continuous con-

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

control tasks, we propose a novel framework called SkillTree. SkillTree introduces the concept of skills, which represent temporal abstractions of low-level actions in the context of RL. These skills simplify the decision-making process by breaking down long trajectories into manageable segments, enabling the agent to plan at a higher level, thereby reducing the complexity associated with extended tasks. However, the challenge of skill representation arises, as skills are often represented in continuous spaces, leading to difficulties in control and explainability. To overcome this, our framework leverages hierarchical structures and discrete skill representation learning. By regularizing the skill space into discrete units, we simplify policy learning for skill selection and enhance the explainability of the learned skills. Specifically, we employ decision trees to index and execute these discrete skills, combining the inherent explainability of decision trees with the flexibility of skill-based representations. This approach provides a robust solution for managing complex, long-horizon decision tasks while offering clear, skill-level insights into the decision-making process. We summarize the main contributions of this paper as follows:

1. We propose SkillTree, a novel hierarchical, skill-based method for explainable reinforcement learning. To the best of our knowledge, it marks the first successful application of DT-based explainable method in long-horizon continuous control tasks.
2. We introduce a method for discrete skill representation learning, which effectively reduces the skill space and improves the efficiency of skill-based policy learning.
3. Experiment results across a variety of robotic manipulation tasks demonstrate that our method provides skill-level explanations while achieving performance comparable to neural network based approaches.

Related Work

Explainable RL

Recently, explainability approaches in DRL have been broadly categorized into intrinsic and post-hoc explanations, based on the method of their generation. We briefly discuss these two classes of methods here. For a more detailed taxonomy and discussion of XRL, please refer to the following surveys: (Heuillet, Couthouis, and Díaz-Rodríguez 2021; Hickling et al. 2023; Milani et al. 2024).

Intrinsic explanation methods are designed with explainability as a fundamental aspect, often incorporating simpler, more transparent models such as decision trees (Liu et al. 2019; Silva et al. 2020; Ding et al. 2020), linear models (Rajeswaran et al. 2017; Molnar, Casalicchio, and Bischl 2020; Wabarth and Pineau 2023) or symbolic rules (Lyu et al. 2019; Ma et al. 2021). By ensuring that the model itself is explainable, intrinsic methods offer real-time, straightforward explanations of the agent’s decisions. For instance, using a decision tree as the policy model enables immediate and clear understanding of the decision-making process.

Post-hoc methods apply explainability techniques after the RL model has been trained. These approaches do not alter the original model but instead seek to explain the

decisions of more complex, often black-box models. Typical post-hoc methods include feature attribution methods like saliency maps (Greydanus et al. 2018; Anderson et al. 2019; Olson et al. 2021) and Shapley value (Zhang et al. 2021; Heuillet, Couthouis, and Díaz-Rodríguez 2022), which highlight the most influential input features. Model distillation approximates the black-box model, either locally or globally, with explainable models like decision trees (Bastani, Pu, and Solar-Lezama 2018; Coppens et al. 2019; Bewley and Lawry 2021; Orfanos and Lelis 2023). Counterfactual explanations generate alternative scenarios by modifying input features to observe changes in the model’s output, offering a way to understand the sensitivity of decisions (Madumal et al. 2020; Olson et al. 2021). Some studies have explored identifying critical states crucial for achieving the final reward from a state-reward perspective (Guo et al. 2021; Cheng et al. 2024). Additionally, researchers have analyzed the impact of different training samples on policy training outcomes (Deshmukh et al. 2023). Lastly, rule extraction methods create human-readable rules that approximate the model’s behavior and enhance the transparency of the decision-making process (Hein, Udluft, and Runkler 2018; Landajuela et al. 2021).

Decision Tree

Decision trees are widely used for their explainability and simplicity, often serving as function approximators in reinforcement learning. Classical decision tree algorithms like CART (Breiman et al. 1984) and C4.5 (Quinlan 1993) produce explainable surrogate policies but are limited in expressiveness and impractical for integration into DRL models. Approaches such as VIPER (Bastani, Pu, and Solar-Lezama 2018) attempt to distill neural network policies into verifiable decision tree by imitation learning. Recent advancements include rule-based node divisions (Dhebar and Deb 2020) and parametric differentiable decision tree, such as soft decision tree (Frosst and Hinton 2017) and differentiable decision tree for approximating Q-function or policy (Silva et al. 2020). While these methods improve expressiveness, they are generally constrained to simpler, low-dimensional environments. By contrast, our approach addresses the complexity of high-dimensional continuous control tasks by converting the action space into a skill space, thereby reducing complexity and providing a more effective and explainable solution for challenging tasks.

Skills in Reinforcement Learning

Recent works have increasingly focused on learning and utilizing skills to improve agent efficiency and generalization (Shu, Xiong, and Socher 2018; Hausman et al. 2018; Shankar and Gupta 2020; Lynch et al. 2020). Skills can be extracted from the offline dataset (Shiarlis et al. 2018; Kipf et al. 2019; Pertsch, Lee, and Lim 2021; Pertsch et al. 2021; Shi, Lim, and Lee 2023) or manually defined (Lee, Yang, and Lim 2019; Dalal, Pathak, and Salakhutdinov 2021; BAAI 2023). Typically, these skills are represented as high-dimensional, continuous latent variables, which limits their explainability. In contrast, we propose extracting a discrete

skill representation from a task-agnostic dataset, making it suitable for DT-based policy learning.

Preliminary

Reinforcement Learning

In this paper, we are concerned with a finite Markov decision process (Sutton and Barto 2018), which can be represented as a tuple $M = (\mathcal{S}, \mathcal{A}, R, \mathcal{P}, p_0, \gamma, H)$. Where \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ is the reward space, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$ is the transfer function, $p_0 : \Delta(\mathcal{S})$ is the initial state distribution, $\gamma \in (0, 1)$ is the discount factor, and H is the episode length. The learning objective is to find the optimal policy π maximizing the expected discounted return

$$\max_{\pi} \mathcal{J}(\pi) = \max_{\pi} \mathbb{E}_{s_0 \sim p_0, (s_0, a_0, \dots, s_H) \sim \pi} \left[\sum_{t=0}^{H-1} \gamma^t \mathcal{R}(s_t, a_t) \right]. \quad (1)$$

Soft Decision Tree

The differentiable decision tree (Frosst and Hinton 2017) is a special type of decision tree that differs from the traditional ones like CART (Breiman et al. 1984) by using probabilistic decision boundaries instead of deterministic boundaries. This modification increases the flexibility of model by allowing for smoother transitions at each decision node, as shown in Figure 1. The differentiable soft decision tree is a complete binary tree in which each inner node can be represented by a learnable weight ω_j^i and bias ϕ_j^i . Here, i and j represent the layer index and the node index of that layer, respectively. By using the sigmoid function $\sigma(\cdot)$, the probability $\sigma(\omega x + \phi)$ represents the transfer probability from that node to the left subtree, and $1 - \sigma(\omega x + \phi)$ gives the transfer probability to the right subtree. This process allows the decision tree to provide a ‘‘soft’’ decision at each node for the current input.

Consider a decision tree with depth d . The nodes in the tree are denoted by n_u , where $u = 2^i - 1 + j$ denotes the node index, and the decision path P can be represented as

$$P = \arg \max_{\{u\}} \prod_{i=0}^{d-1} \prod_{j=0}^{2^i-1} p_{\lfloor \frac{j}{2} \rfloor \rightarrow j}^{i \rightarrow i+1} \quad (2)$$

where $\{u\}$ denotes nodes on the decision path, and $p_{\lfloor \frac{j}{2} \rfloor \rightarrow j}^{i \rightarrow i+1}$ denotes the probability of moving from node $n_{2^i + \lfloor j/2 \rfloor}$ to node $n_{2^{i+1} + j}$, the probability is calculated as

$$p_{\lfloor \frac{j}{2} \rfloor \rightarrow j}^{i \rightarrow i+1} = \begin{cases} \sigma \left(\omega_{\lfloor \frac{j}{2} \rfloor}^i x + \phi_{\lfloor \frac{j}{2} \rfloor}^i \right) & \text{if } j \bmod 2 = 0, \\ 1 - \sigma \left(\omega_{\lfloor \frac{j}{2} \rfloor}^i x + \phi_{\lfloor \frac{j}{2} \rfloor}^i \right) & \text{otherwise.} \end{cases} \quad (3)$$

With Equation 2 and 3, the tree can be traversed down to the leaf nodes. Leaf nodes are represented by parameter vectors $\mathbf{w}^{2^d \times K}$, where K is the number of output categories. The output of the leaf node k is a categorical distribution given by $\text{softmax}(\mathbf{w}_k)$, which is independent of the input.




Figure 1: Comparison of the soft decision tree (left) and the hard decision tree (right).

The learnable parameters of both the leaf nodes and the inner nodes can be optimized using existing DRL algorithms, making the soft decision tree suitable as an explainable alternative model for DRL policy.

Explainable RL with SkillTree

Our goal is to leverage discrete skill representations to learn a skill decision tree, enabling skill-level explainability. Overall, our approach is divided into three stages: (1) extracting the discrete skill embeddings from the offline dataset, (2) training an explainable DT-based skill policy by RL for downstream long-horizon tasks, and (3) distilling the trained policy into a simplified decision tree.

Learning Discrete Skill Embeddings

Our goal is to obtain a fixed number of skill embeddings that contain a temporal abstraction of a sequence of actions. Skills represent action sequences of useful behaviors and are represented by D -dimensional vectors. We assume an existing task-agnostic dataset $\mathcal{D} = \{\tau_i\}$, $\tau_i = \{s_0, a_0, \dots, s_{H_i}, a_{H_i}\}$, consists of d trajectories of varying lengths, where each trajectory includes states s_t, \dots, s_{H_i} and corresponding actions a_t, \dots, a_{H_i} . To regularize the skill embeddings, instead of learning within the continuous skill space, we employ a skill table $Z = \{e_1, e_2, \dots, e_K\}$ (i.e., codebook) that contains D -dimensional vectors, each representing a learnable skill embedding.

To learn reusable skills from offline datasets and guide exploration efficiently, we modified VQ-VAE (Van Den Oord, Vinyals et al. 2017) to learn skill representations and skill prior, as illustrated in Figure 2. First, a demonstration trajectory is randomly divided into state-action segments of fixed length h during the training. Each segment contains h observation states $\mathbf{s}_{t,h} = s_t, s_{t+1}, \dots, s_{t+h-1}$ and the corresponding actions $\mathbf{a}_{t,h} = a_t, a_{t+1}, \dots, a_{t+h-1}$. The input to the encoder q_ϕ is $\mathbf{s}_{t,h}$ and $\mathbf{a}_{t,h}$. The output of the encoder is a D -dimensional embedding z_e . Next, we select the embedding from the codebook that is nearest to z_e and its index $k = \arg \min_j \|z_e - e_j\|_2$, obtaining the regularized embedding $z_q = e_k$. Finally, the low-level actual action a_t is obtained through the state-conditioned decoder $\pi_l(s_t, z_q)$, which serves as the low-level policy.

In addition, following the approach of (Pertsch, Lee, and Lim 2021), we introduce a skill prior to guide the high-level policy. The skill prior $p(k_t | s_t)$ predicts the skill distribution given the first state of the sequence in the offline

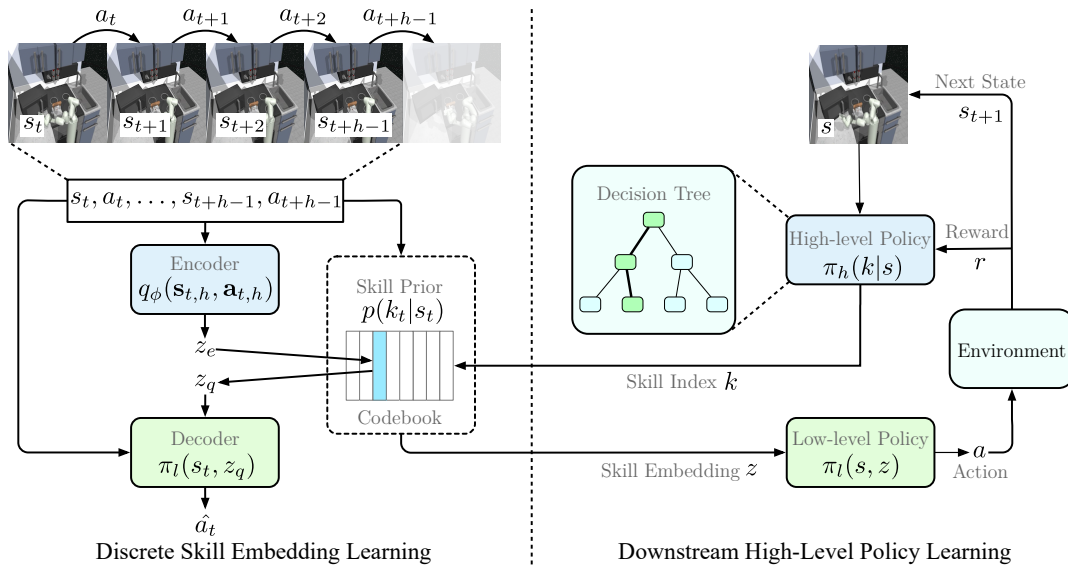


Figure 2: Discrete skill embedding learning and downstream high-level DT policy learning. After completing the skill learning, we freeze the decoder and skill prior, and then proceed to finetune the codebook during the high-level policy learning.

data. This enhances exploration efficiency during high-level policy training by avoiding random skill sampling during warming up. Here, s_t represents the first state of the sequence. The skill prior is a K -categorical distribution that matches the output of encoder. The objective of the skill prior is to fit the output distribution of encoder conditioned solely on the first state s_t . Since the encoder can only produce a fixed skill index k_t , the skill prior aims to match a one-hot distribution of that index, aligning its predictions with the encoder. The overall training objective is

$$\begin{aligned} \mathcal{L} = & \sum_{t=0}^{h-1} \|\pi_l(s_t, z_q) - a_t\|_2^2 \\ & + \|\text{sg}[z_e] - e\|_2^2 + \beta \|z_e - \text{sg}[e]\|_2^2 \\ & - \sum_{k=1}^K q_\phi(z = k | s_{t,h}, \mathbf{a}_{t,h}) \log p(k_t = k | s_t), \end{aligned} \quad (4)$$

where $\text{sg}[\cdot]$ is the stopgradient operator, $p(k_t = k | s_t)$ is the probability that k_t is category k and β is the hyperparameter of the commitment loss term. Categorical distribution $q_\phi(z = k | s_{t,h}, \mathbf{a}_{t,h})$ probabilities are defined as one-hot as

$$q_\phi(z = k | s_{t,h}, \mathbf{a}_{t,h}) = \begin{cases} 1 & \text{for } k = \arg \min_j \|z_e - e_j\|_2, \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where z_e is the output of the encoder.

Downstream RL Policy Learning

To reuse the learning skill embeddings and improve exploration efficiency in RL training, we utilize the skill prior to regularize the high-level policy in skill space. After pre-training the low-level policy and skill prior, the low-level policy $\pi_l(s_t, z_t)$ and the skill prior $p(k_t | s_t)$ are fixed. The

high-level policy $\pi_h(k_t | s_t)$ is executed every h -steps, and the output is the categorical distribution of the skill index k_t . To obtain the actual skill embedding, we sample an index k_t from the distribution and then query the codebook with k_t , i.e., $z_t = Z[k_t]$. In each single time step, the low-level policy predicts the output action conditioned on the state and skill embedding. The low-level policy executes h -steps and the state changes to s_{t+h} . We use $\mathcal{P}^+(s_{t+h} | s_t, z_t)$ to denote the state transition after h steps (see line 3-6 in Algorithm 1). Next, in order to improve the training efficiency by utilizing the prior information obtained from previous skill training, we add an additional prior term as the learning objective:

$$J(\theta) = \mathbb{E}_{\pi_h} \left[\sum_{t=0}^{H-1} \gamma^t r^+(s_t, z_t) - \alpha D_{\text{KL}}(\pi_h(k_t | s_t) \| p(k_t | s_t)) \right], \quad (6)$$

where $r^+(s_t, z_t) = \sum_{i=t}^{t+h} \mathcal{R}(s_i, z_i)$ is the total sum of rewards, θ is the trainable parameter of policy, α is the hyperparameter of divergence term. To prevent the policy from overfitting the prior, following (Pertsch, Lee, and Lim 2021), α can be automatically tuned by defining a target divergence δ (see Algorithm 1 line 15).

Typically, the high-level policy is implemented as a neural network, which is a black box and lacks transparency. In order to achieve an explainable decision making process, we implement the high-level policy $\pi_h(k_t | s_t)$ using a differentiable soft decision tree instead, which can be optimized using existing backpropagation algorithm. Furthermore, the structure of the skill prior is identical to that of the high-level policy, both of which are implemented using soft decision trees. To improve learning efficiency, the parameters of the skill prior are used to initialize the policy during the initialization of training. We modify SAC (Haarnoja et al. 2018) algorithm to learn the objective 6 and the process is summarized in Algorithm 1. See Appendix for more details.

Algorithm 1: SkillTree RL Training

Inputs: codebook Z , target divergence δ , learning rates $\lambda_\theta, \lambda_\psi, \lambda_\alpha$, target update rate τ , DT skill prior $p(k_t|s_t)$ and discount factor γ .

```
1: Initialize replay buffer  $\mathcal{B}$ ,  $d$ -depth high-level DT policy
    $\pi_\theta(k_t|s_t)$ , critic  $Q_\psi(s_t, z_t)$ , target network  $Q_{\bar{\psi}}(s_t, z_t)$ 
2: for each iteration do
3:   for every  $h$  environment steps do
4:      $k_t \sim \pi_\theta(k_t|s_t)$ 
5:      $z_t = Z[k_t]$ 
6:      $s_{t'} \sim \mathcal{P}^+(s_{t+h}|s_t, z_t)$ 
7:      $\mathcal{B} \leftarrow \mathcal{B} \cup \{s_t, z_t, r_t^+, s_{t'}\}$ 
8:   end for
9:   for each gradient step do
10:     $k_{t'} \sim \pi_\theta(k_{t'}|s_{t'})$ 
11:     $z_{t'} = Z[k_{t'}]$ 
12:     $\bar{Q} = r_{t'}^+ + \gamma [Q_{\bar{\psi}}(s_{t'}, z_{t'}) + \alpha D_{\text{KL}}(\pi_\theta(k_{t'}|s_{t'})||p(k_{t'}|s_{t'}))]$ 
13:     $\theta \leftarrow \theta - \lambda_\theta \nabla_\theta [Q_{\bar{\psi}}(s_t, z_t) - D_{\text{KL}}(\pi_\theta(k_t|s_t)||p(k_t|s_t))]$ 
14:     $\psi \leftarrow \psi - \lambda_\psi \nabla_\psi \left[ \frac{1}{2} (Q_\psi(s_t|z_t) - \bar{Q})^2 \right]$ 
15:     $\alpha \leftarrow \alpha - \lambda_\alpha \nabla_\alpha [\alpha (D_{\text{KL}}(\pi_\theta(k_t|s_t)||p(k_t|s_t)) - \delta)]$ 
16:     $\bar{\psi} \leftarrow \tau \psi + (1 - \tau) \bar{\psi}$ 
17:   end for
18: end for
19: return policy  $\pi_\theta(k_t|s_t)$ , finetuned codebook  $Z$ 
```

Distilling the Decision Tree

After downstream RL training, we obtain a soft decision tree with skill-level explainability, where each decision node probabilistically selects child nodes until a skill choice is made. Generally, models with fewer parameters are easier to understand and accept. To further simplify the model structure, we distill the soft decision tree into a more straightforward hard decision tree through discretization. Previous methods select the feature with the highest weight at each node but it often obviously degrades tree performance (Ding et al. 2020). By framing the problem as an imitation learning task, we leverage the learned high-level policy as an expert policy. We sample trajectories from the environment to generate the dataset. Given the simplified skill space, we employ a low-depth CART to classify state-skill index pairs. This approach effectively manages the complexity of the tree while preserving performance, ensuring that the model remains both explainable and efficient.

Experiments

In this section, we evaluate SkillTree’s effectiveness in achieving a balance between explainability and performance. We demonstrate that our method not only provides clear, skill-level explanations but also achieves performance on par with neural network based approaches.

Tasks Setup

To evaluate the performance of our method in long-horizon sparse reward tasks, we chose the *Franka Kitchen* control tasks in D4RL (Fu et al. 2020), robotic arm manipulation




Figure 3: Four long-horizon sparse reward tasks to evaluate. (a) The robotic arm has to finish four subtasks in the correct order, i.e., *Microwave - Kettle - Bottom Burner - Light (MKBL)*. (b) Similar to (a), but with different subtasks: *Microwave - Light - Slide Cabinet - Hinge Cabinet (MLSH)*. (c) Finish subtasks in the correct order, i.e., *Open Drawer - Turn on Lightbulb - Move Slider Left - Turn on LED*. (d) In the office cleaning task, the robotic arm needs to pick up objects and place them in corresponding containers in sequence.

task in *CALVIN* (Mees et al. 2022) and *Office* cleaning task (Pertsch et al. 2021), as illustrated in Figure 3.

Kitchen The *Franka Kitchen* environment is a robotic arm controlled environment based on the mujoco implementation. The environment contains a 7-DOF robotic arm that is able to interact with other objects in the kitchen, such as a kettle and other objects. A total of 7 subtasks are included. The observation is 60-dimensional and contains information about joint positions as well as task parameters, and the action space is 9-dimensional with episode length of 280. The offset dataset contains 601 trajectories, and each trajectory completes a variety of tasks that may not be in the same order, but no more than 3 subtasks. We set two sets of different subtasks: *MKBL* and *MLSH*, as shown in Figure 3(a) and 3(b). The *MLSH* task is more difficult to learn due to the very low frequency of subtask transitions in the dataset.

CALVIN In *CALVIN*, a 7-DOF Franka Emika Panda robotic arm needs to be controlled to open a drawer, light a light bulb, move a slider to the left and light an LED in sequence. the observation space is 21-dimensional and contains both robotic arm states and object states with episode length of 360 steps. The dataset contains 1,239 trajectories




Figure 4: Downstream task learning curves of both our method and baselines. Averaged over 5 independent runs.

Method	Kitchen MKBL		Kitchen MLSH		CALVIN		Office	
	ACS	Leaf	ACS	Leaf	ACS	Leaf	ACS	Leaf
SPiRL	3.08 ± 0.46	N/A	2.82 ± 1.08	N/A	2.80 ± 1.42	N/A	1.48 ± 0.96	N/A
VQ-SPiRL	2.97 ± 0.41	N/A	2.04 ± 0.85	N/A	2.10 ± 0.81	N/A	1.69 ± 0.64	N/A
BC+SAC	0.00 ± 0.00	N/A	0.00 ± 0.00	N/A	0.01 ± 0.10	N/A	0.00 ± 0.00	N/A
CART	2.11 ± 0.92	987	1.16 ± 0.78	981	1.02 ± 0.20	998	0.00 ± 0.00	949
CART+DC	1.32 ± 0.66	970	1.20 ± 0.88	985	1.49 ± 0.50	1006	0.00 ± 0.00	961
SkillTree	2.90 ± 0.42	64	2.58 ± 0.80	64	1.90 ± 0.90	64	2.00 ± 0.00	64
SkillTree (Distillation)	3.06 ± 0.65	64	2.52 ± 0.82	64	2.42 ± 0.49	64	2.00 ± 0.00	64
SkillTree (DC+D)	3.25 ± 0.70	64	2.62 ± 0.76	64	3.00 ± 0.00	64	2.00 ± 0.00	64

Table 1: Average completed subtasks (ACS) and the number of leaf nodes (Leaf) across methods and domains.

and a total of 34 subtasks. Since it contains a large number of subtasks, the average frequency of transitions between any subtasks in the dataset is very low, requiring more precise skill selection, which poses a greater challenge during the learning of the high-level policy.

Office For the *Office* environment, a 5-DOF WidowX robotic arm needs to be controlled to pick up multiple objects and place them in their corresponding containers with episode length of 350. We utilized a dataset collected by (Pertsch et al. 2021), which includes 2,331 trajectories. The state space is 97-dimensional and the action space is 8-dimensional. Due to the freedom of object manipulation in this environment, the task is more challenging to complete.

Baselines

In the experiments, we compare our proposed method against several representative baselines to demonstrate its effectiveness and competitive performance with additional skill-level explainability.

- SPiRL (Pertsch, Lee, and Lim 2021): Learns continuous skills embeddings and a skill prior, and guides a high-level policy using the learned prior. It serves as a competitive skill-based RL algorithm but lacks explainability.
- VQ-SPiRL: Discretizes the skills in SPiRL but uses neural networks as the high-level policy model structure.
- Behavioral Cloning + Soft Actor-Critic (BC+SAC): Trains a supervised behavioral cloning policy from the demonstration dataset and finetunes it on the downstream task using Soft Actor-Critic (Haarnoja et al. 2018). It serves as a general algorithm without leveraging skills.

- CART (Breiman et al. 1984): Imitation learning using CART with trained SPiRL agent as the teacher.
- CART+DC: CART is trained on the data generated by SPiRL after applying data cleaning.

For SkillTree, we set the depth of the tree to 6 in all domains to balance performance and explainability. The size of codebook K is 16 for *Kitchen* and *Office*, and 8 for *CALVIN*. We set skill horizon h to 10 for all experiments. For CART, we sample 1,000 trajectories on each domain to train and set the maximum depth to 10.

Results

From the learning curves shown in Figure 4, it is evident that our method demonstrates comparable learning efficiency and asymptotic performance in all domains. This highlights the effectiveness of our discrete skill representation in handling complex, long-horizon tasks. BC+SAC failed to complete any subtasks because of the sparse reward. In the *CALVIN* and *Office* domain, our method exhibits a marked improvement in early learning efficiency compared to SPiRL and performs comparably to VQ-SPiRL. This advantage underscores the benefit of discrete skills in enhancing exploration efficiency, allowing our approach to quickly identify and leverage useful skills for target task. Through comparing to VQ-SPiRL, we validate that the decision tree can maintain performance while reducing model complexity in simplified discrete skill space.

In Table 1, we evaluate the performance of SkillTree after applying data cleaning (DC) and distillation on a sampled dataset. We consider two sampling methods: directly using




Figure 5: Visualization of the SkillTree (DC+D) with depth 3. $qpos$ and $qpos_obj$ mean the position of the robotic arm and objects, respectively. n denotes the number of state-skill pairs in the divided decision set.

SkillTree for distillation and applying data cleaning before distillation (DC+D). For data cleaning, we retain trajectories that achieve at least two completed subtasks for *Kitchen* and *CALVIN*, and at least one for *Office*. In each environment, we sample 1000 trajectories, ensuring consistency across tasks. To balance performance and explainability, we set the tree depth to 6 across all domains.

The results show that SkillTree DC+D achieves the best performance in most domains, outperforming other baselines. Data cleaning significantly improves performance by filtering low-quality trajectories and emphasizing high-reward ones. For instance, in the DC+D setting, the model demonstrates stronger performance while maintaining simplicity. In contrast, CART+DC also benefits from data cleaning but performs worse than SkillTree, likely due to its limited expressiveness even with excessive number of leaf nodes (~ 1000), which reduce explainability. We also observe that the discrete skill representation (VQ-SPiRL vs. SPiRL) does not significantly enhance performance. This is expected, as the primary goal of discrete skill representation learning is to simplify the skill space, facilitating integration with a decision tree for enhanced policy transparency and explainability, rather than improving raw performance. In comparison, the combination of data cleaning and distillation not only simplifies the SkillTree model but also preserves critical decision-making capabilities, enabling superior performance across diverse tasks.

Explainability

In Figure 5, we visualize the final decision tree obtained after distillation for the *Kitchen MKBL* task. This depth of decision tree is only three but can complete an average of 3.03 subtasks. The very few parameters make the decision-making process clear and easy to understand. Each node makes splits based on one of the observed features. In the




Figure 6: We fix the skill output of high-level policy and evaluate the focus on subtasks of each skill in *Kitchen MKBL* task. We set the number of skills K to 16, as shown in the x-axis. The color means the success rate of each subtask.




Figure 7: Skill index output visualization of an episode in *Kitchen MKBL* task, which finished all 4 subtasks.

60-dimensional observations of the *Kitchen* domain, the last 30 dimensions are task parameters that are independent of the state and thus do not influence decision making. This property is reflected in the learned tree structure because the decision nodes use only the features in the first 30 dimensions (robot arm position and object position) for splitting.

To analyze the effects of each skill, we also evaluated the effectiveness of different skills, as shown in Figure 6. We fixed the skill output individually and evaluated the completion of subtasks. Each skill is capable of completing different subtasks with varying focus. SkillTree selects the appropriate skill based on the observed state, which are then executed by the low-level policy. Figure 7 illustrates the sequence of skill indices chosen in one episode, during which four subtasks were completed. It can be seen that the same skill was repeatedly selected at consecutive time steps to complete specific subtasks, indicating the reuse of particular skills. Moreover, repeated selection of the same skill indicates a relatively long-term and stable skill execution process. By isolating and testing each skill, we can clearly see which skills are most effective for particular subtasks. This makes it easier to understand the role and functionality of each skill within the decision making process. See Appendix for more experiment results as well as visualizations.

Conclusions and Limitations

We proposed a novel hierarchical skill-based explainable reinforcement learning framework designed to tackle the challenges of long-horizon decision making in continuous control tasks. By integrating discrete skill representations and

DT, our method offers a robust solution that enhances explainability without sacrificing performance. The proposed approach effectively regularizes the skill space, simplifies policy learning, and leverages the transparency of the DT. Experiment results demonstrate that our method not only achieves competitive performance compared to neural network based approaches but also provides valuable skill-level explanations. This work represents an important step towards more transparent and explainable RL approaches, paving the way for their broader application in complex real-world scenarios where understanding the decision making process is essential.

Despite the promising results demonstrated, several limitations remain. First, the decision tree high-level policy is not well-suited for image observation inputs. Second, the current approach relies on an offline, task-agnostic dataset for skill learning, which may not always be accessible in practical scenarios. Future work could address these challenges by extending the approach to handle image observations, potentially by incorporating explainable AI techniques into the high-level decision-making process. Additionally, investigating how to adapt SkillTree to multi-agent domains (Yuan et al. 2023) represents an interesting direction for future research.

Acknowledgements

The authors would like to thank the anonymous reviewers for their valuable comments and helpful suggestions. The work is supported by the National Natural Science Foundation of China (62306088) and Natural Science Foundation of Heilongjiang Province (YQ2024007).

References

- Akkaya, I.; Andrychowicz, M.; Chociej, M.; Litwin, M.; McGrew, B.; Petron, A.; Paino, A.; Plappert, M.; Powell, G.; Ribas, R.; et al. 2019. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*.
- Anderson, A.; Dodge, J.; Sadarangani, A.; Juozapaitis, Z.; Newman, E.; Irvine, J.; Chattopadhyay, S.; Fern, A.; and Burnett, M. 2019. Explaining reinforcement learning to mere mortals: an empirical study. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 1328–1334.
- BAAI, P. 2023. Plan4mc: Skill reinforcement learning and planning for open-world minecraft tasks. *arXiv preprint arXiv:2303.16563*.
- Bastani, O.; Pu, Y.; and Solar-Lezama, A. 2018. Verifiable reinforcement learning via policy extraction. In *Advances in Neural Information Processing Systems*, volume 31.
- Bewley, T.; and Lawry, J. 2021. Tripletree: A versatile interpretable representation of black box agents and their environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 11415–11422.
- Breiman, L.; Friedman, J.; Stone, C. J.; and Olshen, R. 1984. *Classification and Regression Trees*. CRC Press.
- Cheng, Z.; Wu, X.; Yu, J.; Sun, W.; Guo, W.; and Xing, X. 2024. Statemask: Explaining deep reinforcement learning through state mask. In *Advances in Neural Information Processing Systems*, volume 36.
- Coppens, Y.; Efthymiadis, K.; Lenaerts, T.; Nowé, A.; Miller, T.; Weber, R.; and Magazzeni, D. 2019. Distilling deep reinforcement learning policies in soft decision trees. In *Proceedings of the IJCAI 2019 workshop on explainable artificial intelligence*, 1–6.
- Costa, V. G.; and Pedreira, C. E. 2023. Recent advances in decision trees: An updated survey. *Artificial Intelligence Review*, 56(5): 4765–4800.
- Dalal, M.; Pathak, D.; and Salakhutdinov, R. R. 2021. Accelerating robotic reinforcement learning via parameterized action primitives. In *Advances in Neural Information Processing Systems*, volume 34, 21847–21859.
- Deshmukh, S. V.; Dasgupta, A.; Krishnamurthy, B.; Jiang, N.; Agarwal, C.; Theocharous, G.; and Subramanian, J. 2023. Explaining RL Decisions with Trajectories. In *International Conference on Learning Representations*.
- Dhebar, Y.; and Deb, K. 2020. Interpretable rule discovery through bilevel optimization of split-rules of nonlinear decision trees for classification problems. *IEEE Transactions on Cybernetics*, 51(11): 5573–5584.
- Ding, Z.; Hernandez-Leal, P.; Ding, G. W.; Li, C.; and Huang, R. 2020. Cdt: Cascading decision trees for explainable reinforcement learning. *arXiv preprint arXiv:2011.07553*.
- Frosst, N.; and Hinton, G. 2017. Distilling a neural network into a soft decision tree. *arXiv preprint arXiv:1711.09784*.
- Fu, J.; Kumar, A.; Nachum, O.; Tucker, G.; and Levine, S. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*.
- Greydanus, S.; Koul, A.; Dodge, J.; and Fern, A. 2018. Visualizing and understanding atari agents. In *International conference on machine learning*, 1792–1801. PMLR.
- Guo, W.; Wu, X.; Khan, U.; and Xing, X. 2021. Edge: Explaining deep reinforcement learning policies. In *Advances in Neural Information Processing Systems*, volume 34, 12222–12236.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1861–1870. PMLR.
- Hausman, K.; Springenberg, J. T.; Wang, Z.; Heess, N.; and Riedmiller, M. 2018. Learning an embedding space for transferable robot skills. In *International Conference on Learning Representations*.
- Hein, D.; Udluft, S.; and Runkler, T. A. 2018. Interpretable policies for reinforcement learning by genetic programming. *Engineering Applications of Artificial Intelligence*, 76: 158–169.
- Heuillet, A.; Couthouis, F.; and Díaz-Rodríguez, N. 2021. Explainability in deep reinforcement learning. *Knowledge-Based Systems*, 214: 106685.

- Heuillet, A.; Couthouis, F.; and Díaz-Rodríguez, N. 2022. Collective explainable AI: Explaining cooperative strategies and agent contribution in multiagent reinforcement learning with shapley values. *IEEE Computational Intelligence Magazine*, 17(1): 59–71.
- Hickling, T.; Zenati, A.; Aouf, N.; and Spencer, P. 2023. Explainability in deep reinforcement learning: A review into current methods and applications. *ACM Computing Surveys*, 56(5): 1–35.
- Jitosho, R.; Lum, T. G. W.; Okamura, A.; and Liu, K. 2023. Reinforcement Learning Enables Real-Time Planning and Control of Agile Maneuvers for Soft Robot Arms. In Tan, J.; Toussaint, M.; and Darvish, K., eds., *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, 1131–1153. PMLR.
- Kipf, T.; Li, Y.; Dai, H.; Zambaldi, V.; Sanchez-Gonzalez, A.; Grefenstette, E.; Kohli, P.; and Battaglia, P. 2019. Compile: Compositional imitation learning and execution. In *International Conference on Machine Learning*, 3418–3428. PMLR.
- Kulhánek, J.; Derner, E.; and Babuška, R. 2021. Visual navigation in real-world indoor environments using end-to-end deep reinforcement learning. *IEEE Robotics and Automation Letters*, 6(3): 4345–4352.
- Landajuela, M.; Petersen, B. K.; Kim, S.; Santiago, C. P.; Glatt, R.; Mundhenk, N.; Pettit, J. F.; and Faissol, D. 2021. Discovering symbolic policies with deep reinforcement learning. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 5979–5989. PMLR.
- Lee, Y.; Yang, J.; and Lim, J. J. 2019. Learning to coordinate manipulation skills via skill behavior diversification. In *International conference on learning representations*.
- Liu, G.; Schulte, O.; Zhu, W.; and Li, Q. 2019. Toward interpretable deep reinforcement learning with linear model u-trees. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2018, Dublin, Ireland, September 10–14, 2018, Proceedings, Part II 18*, 414–429. Springer.
- Liu, G.; Sun, X.; Schulte, O.; and Poupart, P. 2021. Learning tree interpretation from object representation for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 34, 19622–19636.
- Lynch, C.; Khansari, M.; Xiao, T.; Kumar, V.; Tompson, J.; Levine, S.; and Sermanet, P. 2020. Learning latent plans from play. In *Conference on robot learning*, 1113–1132. PMLR.
- Lyu, D.; Yang, F.; Liu, B.; and Gustafson, S. 2019. SDRL: interpretable and data-efficient deep reinforcement learning leveraging symbolic planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 2970–2977.
- Ma, Z.; Zhuang, Y.; Weng, P.; Zhuo, H. H.; Li, D.; Liu, W.; and Hao, J. 2021. Learning symbolic rules for interpretable deep reinforcement learning. *arXiv preprint arXiv:2103.08228*.
- Madumal, P.; Miller, T.; Sonenberg, L.; and Vetere, F. 2020. Explainable reinforcement learning through a causal lens. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 2493–2500.
- Mees, O.; Hermann, L.; Rosete-Beas, E.; and Burgard, W. 2022. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics and Automation Letters*, 7(3): 7327–7334.
- Milani, S.; Topin, N.; Veloso, M.; and Fang, F. 2024. Explainable reinforcement learning: A survey and comparative review. *ACM Computing Surveys*, 56(7): 1–36.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.
- Molnar, C.; Casalicchio, G.; and Bischl, B. 2020. Interpretable machine learning—a brief history, state-of-the-art and challenges. In *Joint European conference on machine learning and knowledge discovery in databases*, 417–431. Springer.
- Olson, M. L.; Khanna, R.; Neal, L.; Li, F.; and Wong, W.-K. 2021. Counterfactual state explanations for reinforcement learning agents via generative deep learning. *Artificial Intelligence*, 295: 103455.
- Orfanos, S.; and Lelis, L. H. 2023. Synthesizing programmatic policies with actor-critic algorithms and relu networks. *arXiv preprint arXiv:2308.02729*.
- Pertsch, K.; Lee, Y.; and Lim, J. 2021. Accelerating reinforcement learning with learned skill priors. In *Conference on robot learning*, 188–204. PMLR.
- Pertsch, K.; Lee, Y.; Wu, Y.; and Lim, J. J. 2021. Demonstration-Guided Reinforcement Learning with Learned Skills. In *5th Annual Conference on Robot Learning*.
- Quinlan, J. R. 1993. C4. 5: Programs for Machine Learning.
- Rajeswaran, A.; Lowrey, K.; Todorov, E. V.; and Kakade, S. M. 2017. Towards generalization and simplicity in continuous control. In *Advances in neural information processing systems*, volume 30.
- Shankar, T.; and Gupta, A. 2020. Learning robot skills with temporal variational inference. In *International Conference on Machine Learning*, 8624–8633. PMLR.
- Shi, L. X.; Lim, J. J.; and Lee, Y. 2023. Skill-based Model-based Reinforcement Learning. In *Conference on Robot Learning*, 2262–2272. PMLR.
- Shiarlis, K.; Wulfmeier, M.; Salter, S.; Whiteson, S.; and Posner, I. 2018. Taco: Learning task decomposition via temporal alignment for control. In *International Conference on Machine Learning*, 4654–4663. PMLR.
- Shu, T.; Xiong, C.; and Socher, R. 2018. Hierarchical and Interpretable Skill Acquisition in Multi-task Reinforcement Learning. In *International Conference on Learning Representations*.

- Silva, A.; Gombolay, M.; Killian, T.; Jimenez, I.; and Son, S.-H. 2020. Optimization methods for interpretable differentiable decision trees applied to reinforcement learning. In *International conference on artificial intelligence and statistics*, 1855–1865. PMLR.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. MIT press.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, volume 30.
- Vasić, M.; Petrović, A.; Wang, K.; Nikolić, M.; Singh, R.; and Khurshid, S. 2022. MoËT: Mixture of Expert Trees and its application to verifiable reinforcement learning. *Neural Networks*, 151: 34–47.
- Wabartha, M.; and Pineau, J. 2023. Piecewise Linear Parametrization of Policies: Towards Interpretable Deep Reinforcement Learning. In *International Conference on Learning Representations*.
- Yuan, L.; Zhang, Z.; Li, L.; Guan, C.; and Yu, Y. 2023. A survey of progress on cooperative multi-agent reinforcement learning in open environment. *arXiv preprint arXiv:2312.01058*.
- Zhang, K.; Zhang, J.; Xu, P.-D.; Gao, T.; and Gao, D. W. 2021. Explainable AI in deep reinforcement learning models for power system emergency control. *IEEE Transactions on Computational Social Systems*, 9(2): 419–427.