

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 FROM OBJECTS TO SKILLS: INTERPRETABLE META-POLICIES FOR NEURAL CONTROL

Anonymous authors

Paper under double-blind review

## ABSTRACT

Despite its success in learning high-performing policies for diverse control and decision-making tasks, deep reinforcement learning remains difficult to interpret and align due to the black-box nature of its neural network representations. Neuro-symbolic approaches improve transparency by incorporating symbolic reasoning, but when applied to low-level actions, they result in overly complex policies. We introduce NEXUS, a hierarchical Reinforcement Learning framework that integrates neural skills with neuro-symbolic meta-policies to balance efficiency and interpretability. In its core, it allows transparent reasoning on disentangled high-level actions (i.e. interpretable skills), greatly reducing complexity of symbolic policies. Object-centric representations enable extracting rewards and meta-policies from language models, while the hierarchical structure allow reasoning over skills rather than atomic actions. We experimentally demonstrate that NEXUS agents are interpretable, less prone to reward hacking, and more robust to [environment simplifications](#). We further evaluate how differing levels of meta-policy interpretability (i.e. purely neural or symbolic) influences performance. Overall, NEXUS enables interpretable and robust control via neuro-symbolic reasoning over high-level skills.

## 1 INTRODUCTION

Recent advancements in Deep Reinforcement Learning have led to highly capable agents on a diverse set of tasks (Mnih et al., 2015; Schulman et al., 2017; Gallici et al., 2024); however, these policies are most often based on neural networks that operate as black boxes and thus exhibit difficult to interpret behaviors that may be misaligned (Rudin, 2019). Without interpretability, identifying misalignment or correcting undesirable behaviors, remains a significant challenge for practitioners (Zahavy et al., 2016; Zhang et al., 2018; Delfosse et al., 2024b).

Neuro-symbolic approaches address this issue by combining neural networks for perceptual grounding with symbolic reasoning modules for decision-making (Delfosse et al., 2023a; Hazra & Raedt, 2023; Acharya et al., 2024; Delfosse et al., 2024b; Luo et al., 2024). These architectures aim to enhance transparency by representing the policies through symbolic structures that are more readily interpretable. Despite their promise, applying symbolic reasoning directly to low-level action spaces often results in policies of prohibitive complexity (cf. Figure 1, [Section B](#)). The combinatorial explosion of symbolic rules at fine-grained action levels undermines interpretability and scalability.

Hierarchical Reinforcement Learning offers an alternative by abstracting sequences of actions into higher-level skills or options (Sutton et al., 1999; Dietterich, 2000; Barto & Mahadevan, 2003). While the framework provides a natural structure, the resulting options, when learned autonomously, are often entangled or have overlapping goals, and the meta-policy remains opaque, both hindering interpretability.

To address these limitations, we draw inspiration from the dual-process theory of cognition (Kahneman, 2011), distinguishing fast, intuitive actions (System 1) from slow, deliberative reasoning (System 2). We emulate this cognitive structure in our hierarchical framework NEXUS, Neural EXecution Under Symbols, which preserves the effectiveness of low-level neural policies (System 1), while combining them with a meta-policy layer of simple, interpretable rule sets (System 2) to maintain clarity about the reasoning behind the active skill in complex environments (cf. Figure 2). We ensure disentangled skills, i.e. each skill corresponds to a distinct situation with clear semantics, by defining their subgoals explicitly through reward functions on objects extracted from the image.

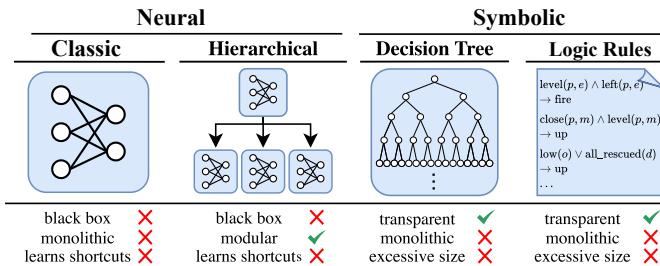


Figure 1: **Current RL policies are not interpretable.** Neural network-based policies are non-modular black box models. Hierarchical RL separates high-level skill selection and low-level control, but the learned entangled skills. Transparent symbolic approaches leverage logical rules or decision trees on symbolic states, but quickly become excessively complex if applied on low-level state features and actions.

To reduce the manual effort required, we leverage the reasoning capabilities of Large Language Models (LLMs) to assist in identifying those skills, defining the high-level meta-policy functions, and generating the reward signals that guide the training of the low-level neural controllers. Both the low-level skills and the meta-policy Q-function are learned jointly in an off-policy manner. Overall, our neuro-symbolic design enables true interpretability on the abstract level of skills without compromising the efficiency of neural agents.

Our primary contributions are as follows:

- (i) We extend *Parallelised Q-Networks* to the hierarchical setting for efficient and scalable meta-policy and skill learning (Section 3.1).
- (ii) We introduce 3 NEXUS variants balancing interpretability and flexibility (Section 3.1 - Section 3.3).
- (iii) We demonstrate that generated object-centric rewards and high-level meta-policy functions guide training towards disentangled skills and interpretable policies (Section 4.2).
- (iv) We provide evidence that NEXUS agents are less susceptible to reward hacking and generalize better to small distribution shift than common algorithms (Section 4.3).

## 2 BACKGROUND

Let us introduce *Deep Reinforcement Learning*, that enables applying neural networks to sequential decision-making tasks. and *Hierarchical Reinforcement Learning*, that decomposes complex tasks into hierarchies of simpler sub-tasks, thereby abstracting actions into skills.

**Deep Reinforcement Learning.** Reinforcement Learning (RL) is a framework for sequential decision making in which an agent learns to interact with an environment in order to maximize cumulative reward. The environment is typically modeled as a Markov Decision Process (MDP), defined by the tuple  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$ , where  $\mathcal{S}$  is the set of states,  $\mathcal{A}$  is the set of actions,  $P(s' | s, a)$  is the transition probability from state  $s$  to state  $s'$  under action  $a$ ,  $R(s, a)$  is the reward function and  $\gamma \in [0, 1]$  is the discount factor. The goal of the agent is to learn a policy  $\pi : \mathcal{S} \rightarrow \mathcal{A}$  that maximizes the expected discounted return:  $\mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$ .

In *Q-learning*, a value-based RL algorithm, the agent seeks to learn the optimal action-value function  $Q^*(s, a)$ , which satisfies the Bellman optimality equation:  $Q^*(s, a) = \mathbb{E}_{s'} [R(s, a) + \gamma \max_{a'} Q^*(s', a')]$ . This function is updated iteratively via the Q-learning update rule:  $Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [r_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)]$ , where  $\alpha$  is the learning rate. With growing state space  $\mathcal{S}$ , it becomes infeasible to store and update a tabular *Q*-function. *Deep Q-Networks (DQN)* (Mnih et al., 2015) address this by using a deep neural network  $Q_\theta(s, a)$ , parameterized by  $\theta$ , to approximate  $Q(s, a)$ . DQN introduces several key modifications to stabilize learning, including experience replay, where transitions  $(s_t, a_t, r_t, s_{t+1})$  are stored in a replay buffer and sampled randomly to break correlations between consecutive updates and target networks, where a separate network  $Q_{\theta-}(s, a)$  is used to compute the target value and is updated periodically:  $y_t = r_t + \gamma \max_{a'} Q_{\theta-}(s_{t+1}, a')$ .

More recently, Gallici et al. (2024) introduced *Parallelised Q-Networks (PQN)*, a simplified variant of DQN that eliminates the use of experience replay and target networks. Instead, PQN leverages a large number of parallel (ideally vectorized) environments and applies normalization techniques to mitigate training instabilities. This high degree of parallelization enables substantially faster training while maintaining performance comparable to state-of-the-art RL algorithms.

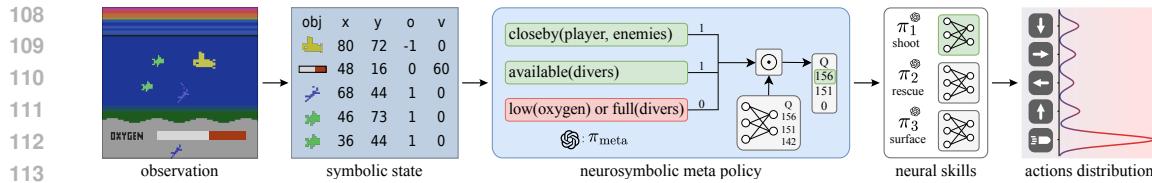


Figure 2: **NEXUS with a neuro-symbolic meta-policy.** The symbolic state is extracted from the observation. The LLM-generated meta policy selects potential relevant skills, then selected using a trained meta Q-function. Finally, the selected neural skill outputs the final low-level action. While the meta policy’s Q-function is trained using the environmental reward, each skill is trained using its own LLM-generated reward signal, allowing for interpretable disentangled skills.

**Hierarchical Reinforcement Learning.** Hierarchical reinforcement learning (HRL) decomposes complex tasks into simpler sub-tasks, often modeled as temporally extended actions or "options" (Sutton et al., 1999). At inference time, a meta policy selects the option, enabling more efficient exploration and transfer in environments with long horizons or sparse rewards. Sub-policies can be learned autonomously using intrinsic objectives (Bacon et al., 2017; Vezhnevets et al., 2017) or guided by manually designed reward functions (Sutton et al., 1999; Dietterich, 2000). However, autonomous discovery often suffers from option entanglement or collapse. NEXUS instead enforces disentangled options via explicit reward functions, which are LLM-generated.

### 3 NEURAL EXECUTION UNDER SYMBOLS

We address the challenge of building interpretable and high performing reinforcement learning agents with NEXUS, a hierarchical framework that combines symbolic reasoning with neural skill execution. The overall pipeline is described in Figure 2. **At a high level, NEXUS decomposes decision-making into two layers: (1) a high-level meta-policy and (2) a set of low-level neural skill policies optimized for specific sub-goals.** At each step, the decision process proceeds as follows. First, object-centric representations are extracted from pixel-based inputs. Second, symbolic rules, encoding the activation conditions of each skill, filter out skills without valid preconditions. Third, the learned meta-policy Q-function selects the most suitable skill among the remaining candidates. Finally, the chosen skill executes its policy to maximize the corresponding sub-goal. Object-centric representations are crucial for enabling symbolic rule definition and skill-specific reward design, both of which are difficult without such structured representations. Moreover, the explicit availability of object information allows the use of LLMs for generating the rules and reward functions.

In the following, we describe three distinct meta-policy function type: purely neural, purely symbolic, and neuro-symbolic. The first method, NEXUS (neural), corresponds to a hierarchical PQN method, which learns both a neural meta-policy and neural skills (based on skill-specific rewards) and which serves as the foundation for subsequent variants. While it offers high flexibility due to the learned meta-policy, it is missing reasoning for why a given sub-goal (and its corresponding skill) is selected in the current step. This shortcoming can be facilitated by replacing the learned meta-policy with a fixed, interpretable function that can be either manually defined or LLM-generated. This variant, NEXUS (symbolic), is fully symbolic at the meta-policy level, prioritizing transparency over adaptability. Finally, we present the hybrid approach NEXUS (nesy) that filters candidate skills via predefined symbolic rules and selects among them using Q-learning. This last variant results in a neuro-symbolic meta-policy, symbolic in rule-based filtering and neural in value estimation, combined with neural sub-policies, achieving a principled trade-off between interpretability and flexibility.

#### 3.1 HIERARCHICAL PQN

We accommodate the hierarchical learning setup by modeling the environment as two layers of MDPs. On the action level, we define a collection of MDPs  $\mathcal{M} = \{\langle \mathcal{S}, \mathcal{A}, P, R_n, \gamma \rangle\}_{n=1}^N$  for  $N$  skills, each associated with a distinct option policy  $\pi_n \in \Pi$  that decides over the actual environment actions. The meta-level MDP  $\mathcal{M}_{\text{meta}} = \langle \mathcal{S}, \Pi, P, R_{\text{env}}, \gamma \rangle$  governs option selection through a meta-policy  $\pi_{\text{meta}}$  that selects over the options  $\Pi$ . Note that the option-level MDPs differ only in their reward functions  $R_n$ , while the meta-level MDP retains the environment’s original reward  $R_{\text{env}}$ .

162 To enable learning across sub-policies even when they are not actively selected, we adopt off-policy  
 163 Q-learning. This allows all the skills to learn from trajectories generated by the active skill and  
 164 meta-policy while optimizing for their respective reward structures. Exploration within each skill  
 165 is conducted via  $\epsilon$ -greedy action selection. For scalability through vectorization, we build our  
 166 hierarchical training approach upon PQL with  $\lambda$ -returns (Gallici et al., 2024). However, instead  
 167 of learning a single shared Q-network  $Q_\phi$  as in PQL, we learn separate Q-functions  $Q_{\phi_n}$  for each  
 168 low-level skill, along with a meta-level Q-network  $Q_{\phi_{\text{meta}}}$ . Each Q-network thus predicts the returns  
 169 of its own reward function.

170 For updating, we roll out trajectories  $(s_0, \dots, s_T)$  following the global policy by selecting the next  
 171 active skill  $\pi_n \in \Pi$  using the meta-policy Q-function  $Q_{\phi_{\text{meta}}}$  and find the most promising action  
 172 according to the active skill’s Q-function  $Q_{\phi_n}$ :

$$\pi_n = \arg \max_{\pi_n \in \Pi} Q_{\phi_{\text{meta}}}(s_t, \pi_n), \quad a' = \arg \max_{a' \in \mathcal{A}} Q_{\phi_n}(s_t, a'). \quad (1)$$

173 By executing the action in the environment, we obtain the rewards and the next environmental state for  
 174 the next iteration. We employ  $\epsilon$ -greedy for exploration during both the skill and the action selection.

175 Next to the environment reward  $r_{\text{env},t}$ , we require skill-specific reward functions  $r_{n,t}$  that are based  
 176 on the object-centric state  $s_t$  and may be automatically generated by LLMs (cf. Section G), which  
 177 we use to compute the  $\lambda$ -returns recursively back in time (for details, we refer the reader to Gallici  
 178 et al. (2024); Daley & Amato (2019)):

$$R_{n,t}^\lambda = r_{n,t} + \gamma \left[ \lambda R_{n,t+1}^\lambda + (1 - \lambda) \max_{a'} Q_{\phi_n}(s_{t+1}, a') \right], \quad (2)$$

183 and similarly, for the learned meta-policy using environment rewards:

$$R_{\text{env},t}^\lambda = r_{\text{env},t} + \gamma \left[ \lambda R_{\text{env},t+1}^\lambda + (1 - \lambda) \max_{\pi_n} Q_{\phi_{\text{meta}}}(s_{t+1}, \pi_n) \right], \quad (3)$$

186 or, if  $s_t$  is terminal  $R_{n,t}^\lambda = r_{n,t}$  and  $R_{\text{env},t}^\lambda = r_{\text{env},t}$ . All learned Q-functions are updated towards their  
 187  $\lambda$ -returns. We provide the full algorithm in Section C.

### 189 3.2 INTERPRETABLE META-POLICY FUNCTION

191 While the hierarchical structure allows to identify the active skill, it does not expose the decision  
 192 process that leads to its selection. To make this process transparent, we introduce interpretable  
 193 meta-policy functions implemented as rule-based programs. We choose this representation for two  
 194 key reasons. First, rule-based programs offer simplicity and accessibility since conditional statements  
 195 can typically be understood and modified even by users with limited programming background.  
 196 Second, LLMs are highly effective at generating and editing code-like structures due to extensive  
 197 pretraining on programming data. This enables generation of human-interpretable meta-policies and  
 198 greatly reduces manual effort. Further details on how we employ LLMs to generate the meta-policy  
 199 function are available in Section G.

200 Formally, instead of relying on a learned Q-function  $Q_{\phi_{\text{meta}}}$ , we define a meta-policy  $\pi_{\text{meta}} : \mathcal{S} \rightarrow \Pi$   
 201 as a set of human-readable rules that directly maps the current object-centric state  $s_t \in \mathcal{S}$  to the  
 202 selected low-level policy  $\pi_n \in \Pi$ :

$$\pi_n = \pi_{\text{meta}}(s_t), \quad (4)$$

204 We hereby constrain  $\pi_{\text{meta}}$  to simple rules that are mutually exclusive to enable transparent inspection  
 205 of policy decisions. The selected skill can be traced back to the specific rule that evaluates to true  
 206 given the current state (cf. Figure 5, left side).

### 208 3.3 NEURO-SYMBOLIC META-POLICY FUNCTION

210 The symbolic meta-policy is a deterministic function that selects skills based on the rule with the  
 211 highest priority. However, this approach may become inefficient in situations where multiple skills  
 212 are simultaneously applicable. Consider an agent controlling a submarine that can shoot a nearby  
 213 enemy or surface to replenish oxygen. The most viable depends on the specific context: if the oxygen  
 214 level is critically low, resurfacing may be prioritized; if an enemy is dangerously close, attacking may  
 215 take precedence. Crafting hard-coded rules to handle such trade-offs would quickly lead to growing,  
 less interpretable policies.

Instead, we propose a neuro-symbolic (NeSy) approach, by maintaining a set of high-level, interpretable conditions (e.g., “go to surface if oxygen is low”, “fight if enemy is close”), represented as a binary condition vector  $c_t \in \{0, 1\}^N$ , where each entry indicates whether a particular condition is active at time  $t$ . We formulate these rule-sets again as simple code functions and leverage LLMs for their generation (cf. Figure 5, right side). However, this time we allow multiple conditions to be active concurrently. The skill selection is then modulated by these conditions using a binary mask applied to the meta-policy Q-values:

$$\pi_n = \arg \max_{\pi_n \in \Pi} (c_t \odot Q_{\phi_{\text{meta}}}(s_t, \pi_n)) \quad (5)$$

where  $\odot$  denotes element-wise multiplication,  $Q_{\phi_{\text{meta}}}(s_t, \cdot) \in \mathbb{R}^N$  is the vector of meta-policy Q-values for each interpretable condition, and  $\Pi$  is the set of available skills.

This way, we can preserve interpretability at the symbolic level while enabling the agent to resolve ambiguous scenarios adaptively based on learned preferences. An overview of the different meta-policy variations is provided in Algorithm 1.

**Prior knowledge requirements.** There are two components of NEXUS that require prior knowledge. The first is the specification of skill reward functions. The second is the definition of meta-policy rules for the symbolic and the neuro-symbolic variants. Although these elements can be provided manually, we leverage LLMs to minimize expert input. Given the prompts in Section G, only the game manual is ultimately required.

## 4 EXPERIMENTAL EVALUATION

This work aims to develop an object-centric pipeline for agents capable of solving complex environments while maintaining interpretability of high-level goals. We assess this objective by addressing the following research questions:

- (Q1) Does NEXUS learn meaningful disentangled skills?
- (Q2) Are NEXUS policies interpretable?
- (Q3) Can NEXUS compete with other deep methods?
- (Q4) Does NEXUS improve robustness to game simplifications?
- (Q5) How does noise in the object extractor influence NEXUS?

### 4.1 EXPERIMENTAL SETUP

Our experiments assume access to pre-extracted object representations and attributes. We therefore provide the agent with symbolic environment states directly, allowing us to remain agnostic to the specific object extraction pipeline. We conduct experiments on *JAXAtari*<sup>1</sup>, a JAX-based reimplementation of the Atari Learning Environment (ALE) (Bellemare et al., 2013), specifically using the games *Kangaroo* and *Seaquest*, with object-centric representations similar to *OCatari* (Delfosse et al., 2024a). These games are selected because their gameplay naturally decomposes into low-level skills that must be combined to solve the overall task. In Kangaroo, the objective is to ascend through platforms while avoiding enemies and collecting berries. In Seaquest, the player navigates a submarine to rescue divers and return them to the surface, while avoiding hazards such as sharks

<sup>1</sup><https://github.com/k4ntz/JAXAtari>

---

### Algorithm 1 NEXUS Variants

---

**Require:** Skill policies  $\{\pi_n\}_{n=1}^N$  with Q-functions  $\{Q_{\phi_n}\}_{n=1}^N$ , meta-policy  $Q_{\phi_{\text{meta}}}$  (for learned/soft)

- 1: **for** each episode **do**
- 2:   **for** each environment step (in parallel) **do**
- 3:     Extract object-centric state  $s_t$
- 4:     Select skill  $\pi_{n,t}$  as follows:
  - A) Neural:  $\pi_{n,t} = \arg \max_{\pi_n} Q_{\phi_{\text{meta}}}(s_t, \pi_n)$
  - B) Symbolic:  $\pi_{n,t} = \pi_{\text{meta}}(s_t)$
  - C) NeSy:  $\pi_{n,t} = \arg \max_{\pi_n} (c_t \odot Q_{\phi_{\text{meta}}}(s_t, \pi_n))$
- 5:     Select action  $a_t$  using  $\epsilon$ -greedy from  $\pi_{n,t}$
- 6:     Execute  $a_t$  and observe transition  $(\{r_{n,t}\}_{n=0}^N, r_{\text{env},t}, s_{t+1})$
- 7:   **end for**
- 8:   **for** each gradient step **do**
- 9:     Compute  $R_{n,t}^\lambda$  and  $R_{\text{env},t}^\lambda$  (if A or C)
- 10:    Update each  $Q_{\phi_n}$  using skill-specific  $R_{n,t}^\lambda$
- 11:    Update  $Q_{\phi_{\text{meta}}}$  using  $R_{\text{env},t}^\lambda$  (if A or C)
- 12:   **end for**
- 13: **end for**

---

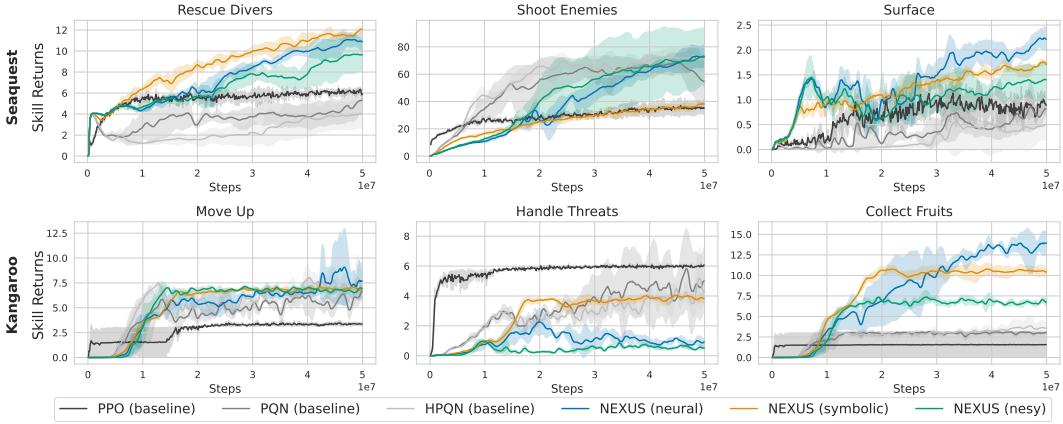


Figure 3: **NEXUS learns disentangled skills from off-policy data.** In Seaquest (top), the baseline methods mainly focus on shooting enemies, while the NEXUS approaches acquire the target skills more evenly. Similarly, in Kangaroo (bottom), the NEXUS approaches learn ‘Move Up’ and ‘Collect Fruits’ reliably, while the baselines focus mostly on ‘Handle Threats’.

and enemy submarines. These two games are known to exhibit *reward hacking* behavior, meaning that the agents find simple shortcuts of increasing their return rather than the intended way that humans would typically choose. Examples include exclusively shooting enemies in Seaquest and only catching falling apples or boxing monkeys in Kangaroo. Additionally, we evaluate on *Crafter* (Hafner, 2021), using the JAX reimplementation by Matthews et al. (2024), an environment that integrates elements from the games Minecraft and NetHack. Unlike Atari games, Crafter presents a more complex, open-ended challenge that requires agents to acquire and coordinate multiple skills, including navigation, combat, resource collection, and crafting. While primarily designed as a benchmark for open-ended learning, Crafter’s task diversity makes it well-suited for evaluating the generalization and skill composition capabilities of RL methods.

Finally, for generalization we also test the pre-trained agents on slightly modified versions of the games. These modifications are designed to highlight the sensitivity of standard reinforcement learning agents to minor changes in the environment, including simplifications. For example, removing enemies from Seaquest to reduce the task’s complexity already leads to a substantial drop in many deep agents performances (Delfosse et al., 2025).

We compare all three variations of NEXUS to the default, non-interpretable PQN (Gallici et al., 2024), and, in case of the Atari games, with the actor-critic PPO (Schulman et al., 2017) as baselines. Additionally, we compare to HPQN, a hierarchical PQN variation that does not employ skill-specific rewards and has a purely neural meta-policy. All results (including baselines) are directly trained on object-centric inputs, which may affect performance compared to image based training. We adhere to the standard frame budgets of both games: 200M frames for JAXAtari and 1B frames for Crafter. Further implementation details are available in Section D.

## 4.2 INTERPRETABILITY RESULTS

**Disentangled skill learning (Q1).** We first evaluate whether NEXUS enables efficient learning of meaningful skills. Since each skill requires its own reward function, we leverage the reasoning capabilities of LLMs to generate them based on the game manual, skill definitions, and the object-centric state (cf. Section G for further information).

Figure 3 demonstrates that NEXUS is able to learn most target skills successfully from a single source of off-policy data. This means that a skill can learn from another skills action and is not required to be activated to do so. Unlike the baselines, which tend to mostly focus on a single skill that maximizes environment reward (e.g., Shoot Enemies for Seaquest, Handle Threats for Kangaroo), NEXUS approaches promote balanced skill acquisition. Especially the symbolic meta-policy approach seems to learn all skills most reliably.

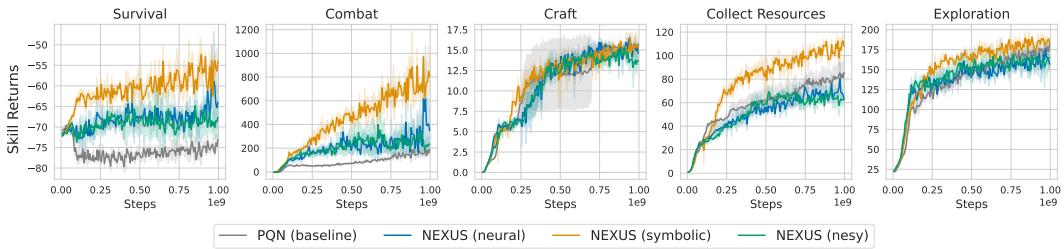


Figure 4: **NEXUS successfully learns LLM-proposed Crafter skills.** Compared to the baseline, NEXUS approaches often converge faster on the proposed skills.

```

336 def meta_policy(st: state):
337     if enemy_close(st.enemies,
338         st.player):
339         return fight_enemies()
340     elif is_available(st.divers):
341         return rescue_divers()
342     elif is_low(st.oxygen):
343         return surface()
344     elif all_collected(st.divers):
345         return surface()
346     return rescue_divers()
347
348 def meta_policy_rules(st: state):
349     fight_enemies = False
350     rescue_divers = True
351     surface = False
352     if enemy_close(st.enemies,
353         st.player):
354         fight_enemies = True
355     if is_low(st.oxygen) or
356     all_collected(st.divers):
357         surface = True
358     return [fight_enemies,
359             rescue_divers, surface]

```

Figure 5: **NEXUS policies and rules are clear and interpretable.** A symbolic meta-policy (left) and similar filtering rules for the neuro-symbolic meta-policy (right) for the Atari game Seaquest.

We experience a similar picture when evaluating on Crafter. Using the game manual and state description, we query an LLM for essential skills. Recurrent versions of NEXUS are then trained explicitly for these skills and compared to a recurrent PQN baseline. Results are presented in Figure 4, where we compare the ability to learn the skills between the different methods. Similar to before, NEXUS approaches are able to learn the five skills with symbolic being the fastest.

**Interpretability of NEXUS policies (Q2).** To assess the interpretability of NEXUS, we visualize a fixed meta-policy for Seaquest on the left side in Figure 5. By abstracting raw observations into object-centric representations and low-level actions into high-level skills, the decision-making process becomes transparent. Each option has a clear, mutually exclusive activation condition, enabling unambiguous skill selection. For example, the combat skill activates if and only if an enemy is in close proximity. This simplicity offers two advantages: technical users can design fully interpretable policies, and LLMs can autonomously generate such meta-policies, which remain editable due to their transparency. We illustrate LLM-based meta-policy generation in Section G.

Full transparency is not guaranteed for neuro-symbolic meta-policies, as multiple conditions may be simultaneously satisfied (e.g., an enemy is nearby and oxygen is low; cf. Figure 5, right side). Nonetheless, such overlaps are rare, and interpretability is largely preserved. In ambiguous cases, interpretability is traded for flexibility, allowing the agent to select the skill with the highest expected return based on the meta-policy Q-values. An example of a Seaquest agent operating under a neuro-symbolic meta-policy is available in Figure 6.

### 4.3 PERFORMANCE RESULTS

**Comparison to other approaches (Q3).** We evaluate NEXUS against baselines using two metrics: Game returns and aligned environment goals, which track progress toward goals defined in game manuals (e.g., divers rescued in Seaquest, level completion in Kangaroo). The latter aims to capture the overall alignment with the intended game objectives, which differ from just maximizing the reward for these environments, where deep agents usually perform reward hacking (Shihab et al., 2025).

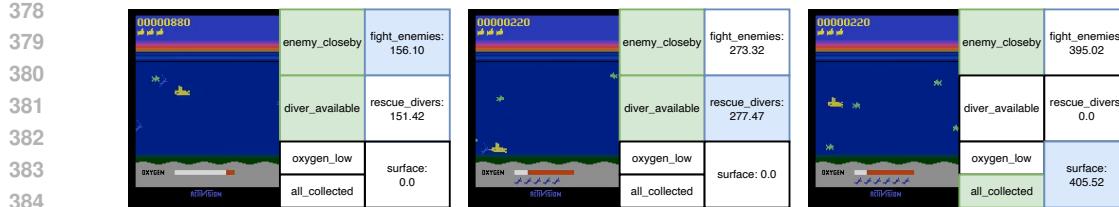


Figure 6: **NEXUS produces interpretable yet flexible high-level plans for ambiguous scenarios.** Left: Both "enemy\_closeby" and "diver\_available" rules evaluate to true; the learned meta-policy prioritizes fighting, likely due to the diver's proximity to the enemy. Middle: Under similar conditions, rescuing the diver is preferred, reason could be the easier access. Right: With all divers collected and an enemy nearby, the meta-policy opts to return to the surface due to the higher estimated return.

As shown in Figure 7, NEXUS approaches are competitive to the baselines. In most cases, they achieve comparable or even higher HNS and notably outperform the baselines on the actual games' main goal, demonstrating the baselines reward hacking tendency. The results suggest that NEXUS mitigates reward misalignment by incorporating domain priors into the decision-making process, while still achieving good performance. [We provide further comparisons to neuro-symbolic and interpretable RL methods in Section E.](#)

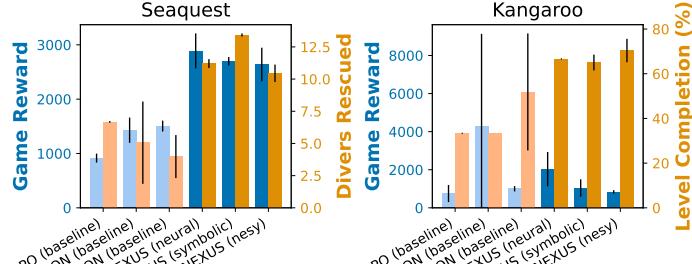


Figure 7: **NEXUS is competitive.** In Seaquest, NEXUS outperforms the baselines on both HNS and rescuing divers. In Kangaroo, NEXUS approaches are better aligned to the actual game objective.

**Robustness to game simplifications (Q4).** Most RL algorithms struggle to adapt to even minor variations in the environment (Delfosse et al., 2025). Surprisingly, their performances drop even in settings that simplify the game for humans, such as removing deadly threats like enemies and their projectiles in Seaquest and Kangaroo. We assess whether NEXUS agents can generalize to such simplified variants by training the agents on the standard versions of the games and evaluating on the unseen simplifications. In Seaquest and Kangaroo we remove all threats, while in Crafter we remove the need to drink water for survival. Figure 8 presents the results. As expected, the baselines suffer substantial performance degradation under the simplifications in all three games.

All NEXUS variants exhibit distinct robustness characteristics. In Atari, only the symbolic variant demonstrates improved robustness, with smaller performance drops in Kangaroo and even gains in Seaquest, while the fully neural and neuro-symbolic variants show limited robustness. In Crafter all NEXUS variants maintain or improve performance. We attribute this robustness to symbolic steering of the meta-policy, which deactivates unnecessary skills such as handling enemies or drinking water.

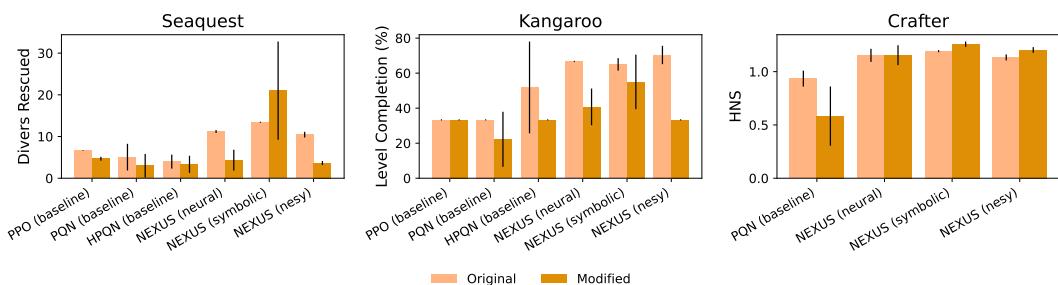


Figure 8: **Symbolic NEXUS remains practical on simplified games.** While baseline performance drops significantly, neuro-symbolic and fully symbolic NEXUS improve in the modified Seaquest (left) and symbolic NEXUS shows smaller performance drops in modified Kangaroo (middle).

432  
 433 **Influence of noisy detections**  
 434 **on NEXUS (Q5).** While ob-  
 435 ject detection methods become  
 436 increasingly reliable, misdetec-  
 437 tions still happen. With the fol-  
 438 lowing ablation, we test whether  
 439 NEXUS is robust to misdetec-  
 440 tions and noise in the detec-  
 441 tions. For that, we incorporate  
 442 a 10% misdetection chance and  
 443 add gaussian noise with a stan-  
 444 dard deviation of 3px to each ob-  
 445 ject attribute during training. The  
 446 results are visualized in Figure 9.  
 447 We observe that both the neu-  
 448 ral and the symbolic approaches  
 449 lose performance when the noise is  
 450 applied in Seaquest, while the  
 451 neuro-symbolic meta-policy  
 452 remains reliable. Surprisingly,  
 453 the experiments on Kangaroo  
 454 indicate that all approaches  
 455 increase the level completion  
 456 score and often times the game  
 457 reward. Most notably, the  
 458 neuro-symbolic meta-policy  
 459 is able to increase the level  
 460 completion rate from  $\sim 70\%$  to  
 461 above 125% and always  
 462 finishes level one. Further details  
 463 and results are available in Section F.

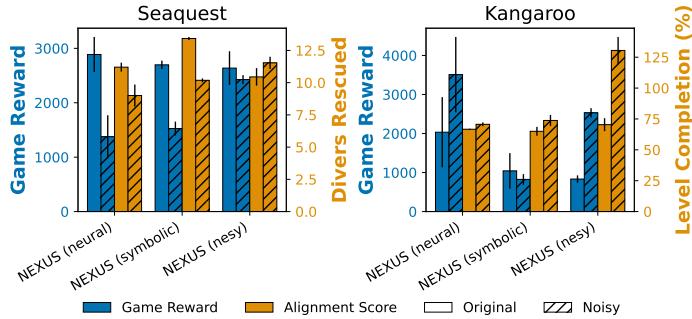


Figure 9: **Effect of noisy detections.** In Seaquest, the neural and the symbolic meta-policy take a performance hit in both reward and alignment score, while the NeSy variation is robust. In Kangaroo, all approaches improve with the added noise.

454 **On the choice of the meta-policy.** Drawing from our findings, we can now offer recommendations  
 455 regarding the optimal design of the meta-policy. The purely neural meta-policy presents the simplest  
 456 training paradigm as it does not need a separate symbolic meta-policy, and it demonstrates strong per-  
 457 formance in terms of training environment reward. However, this approach sacrifices interpretability  
 458 and also performance when evaluated on simplified environments. In contrast, the purely symbolic  
 459 meta-policy necessitates an additional step of rule definition. Since this can be largely automated  
 460 with LLMs, this investment is often justified by its enhanced interpretability and robustness. Lastly,  
 461 the neuro-symbolic approach eases the definition of the additional rules, since they do not need to be  
 462 mutually exclusive. While the performance is often similar to the symbolic policy, it is less robust to  
 463 game simplification, but more robust to noise in the detection method. Considering these trade-offs,  
 464 we advocate for both, the purely symbolic and the neuro-symbolic meta-policies as effective choices,  
 465 offering a compelling balance of strong performance, interpretability, and generalizability.

466  
 467 **Limitations.** NEXUS relies on object-centric scene decoders that accurately provide the agent with  
 468 objects and their positions and sizes from raw images, which we assume to exist in this work. For  
 469 Atari games, multiple approaches are viable (Li et al., 2017; Locatello et al., 2020; Lin et al., 2020;  
 470 Delfosse et al., 2023b) with some achieving near-perfect sprite extraction (Smirnov et al., 2021).  
 471 For real-world data, recent advancements have significantly improved robustness, with models like  
 472 SAM2 (Ravi et al., 2025) reaching up to 90%  $\mathcal{J}$  &  $\mathcal{F}$  accuracy on zero-shot segmentation.

473 NEXUS requires a pre-defined description of the task, i.e. the game’s manual, for the reward  
 474 generation (as done in Wu et al. (2023)). This limits its applicability to tasks that can be explained  
 475 in language. Moreover, while the LLM-generated reward functions and meta-policies generally  
 476 capture valid semantics and are logically consistent, some manual adjustments to align them with the  
 477 implementation framework are still necessary. Prior work on LLM-based reward design (e.g., Xie et al.  
 478 (2024), Ma et al. (2024), Kaufmann et al. (2024)) has documented recurring issues such as reward  
 479 misspecification, proxy objective selection, and over- or under-constrained preconditions, highlighting  
 480 the need for careful verification. NEXUS currently addresses these risks through manual inspection of  
 481 the generated code, but does not incorporate automated diagnostics or iterative improvements during  
 482 learning. Extending our framework by incorporating more refined mechanisms is a crucial direction  
 483 for enhancing robustness and scalability. Lastly, the presented approach is based on Q-learning and  
 484 thus currently limited to discrete action space, however, extension to continuous action spaces seems  
 485 viable by adopting an off-policy actor critic instead of  $\epsilon$ -greedy action selection (Lillicrap et al.,  
 2016; Haarnoja et al., 2018).

486 5 RELATED WORK  
487  
488

489 **Interpretable and Hierarchical Reinforcement Learning.** Interpretability in RL can be introduced  
490 at various stages of the pipeline (Glanois et al., 2024), often by deriving symbolic state representations  
491 from raw observations via object recognition or segmentation (Li et al., 2017; Locatello et al., 2020;  
492 Kirillov et al., 2023; Lin et al., 2020; Delfosse et al., 2023b). Such object-centric states have enabled  
493 interpretable policies through decision trees (Silva et al., 2019; Likmeta et al., 2020; Delfosse et al.,  
494 2024b), logic rules (Maes et al., 2012; Akroud et al., 2019; Delfosse et al., 2023a), **parametric**  
495 **functions (Luo et al., 2024)** or programmatic policies and trees (Verma et al., 2019; Anderson et al.,  
496 2020; Kohler et al., 2024). Neural and logical policies can also be efficiently combined (Shindo et al.,  
497 2025). Complementary efforts introduce hierarchical decompositions, where high-level interpretable  
498 policies select among low-level sub-policies, leveraging annotated task sketches (Andreas et al.,  
499 2017) or (differentiable) symbolic planning (Leonetti et al., 2016; Yang et al., 2018; Jin et al.,  
500 2022; Lyu et al., 2019; Ye et al., 2025). Hierarchical symbolic planning based on object-centric  
501 representations has been shown to be beneficial for task transferability (James et al., 2022) and  
502 robotics applications (Sharma et al., 2020). Our work extends these directions by enforcing semantic  
503 separation of skills, introducing a neurosymbolic meta-policy to balance interpretability and flexibility,  
504 and integrating LLMs throughout the pipeline. Unlike prior approaches, we validate on challenging  
505 Atari and Crafter environments.

506 **Relation to the Options framework.** HRL has been studied extensively, with the Options framework  
507 (Sutton et al., 1999) as the most prominent formulation. NEXUS instantiates the Options framework  
508 by performing intra-option learning, where multiple options are learned simultaneously from shared  
509 off-policy experience, using sub-policies guided by option-specific reward functions and coordinated  
510 by a neuro-symbolic meta-policy. Key differences arise in temporal abstraction and in the treatment  
511 of initiation and termination conditions. Rather than executing an option until termination, the  
512 meta-policy selects the active option at every time step, jointly determining activation and termination.  
513 The logical rule set used to filter Q-values before selection serves as a generalized initiation set,  
514 enabled by the object-centric encoding of observations. **NEXUS’ key innovations to HRL are: (1)**  
515 **disentangled neural options through specialized rewards and (2) interpretable meta-policies.**

516  
517 6 CONCLUSION  
518  
519

520 In this work, we present NEXUS, a hierarchical RL method that combines high interpretability on  
521 meta-policy level with neural, low-level action execution. The evaluation suggest several advantages  
522 of our approach. It learns disentangled sub-policies corresponding to identifiable skills, provides  
523 interpretable and modular structures for inspection and intervention, reduces reward-hacking through  
524 fine-grained control, and remains robust to environment simplifications where standard deep RL  
525 agents fail. We also demonstrate that LLMs can be integrated into the NEXUS pipeline to generate  
526 skill decompositions, reward functions, and symbolic meta-policies, enabling dynamic adaptation to  
527 novel objects, evolving environments, or shifting task objectives. Future work should increase the  
528 autonomy of policy adaptation by incorporating mechanisms to detect when new skills or meta-policy  
529 updates are required, e.g. via causal world models (Yang et al., 2025; Dillies et al., 2025), which could  
530 enable scaling to complex, open-ended environments. **Updates to the skills or meta-policy could be**  
531 **retrieved by re-querying an LLM. Additionally, future work should evaluate the actual interpretability**  
532 **of NEXUS in a user-study.** Overall, this work advances RL interpretability and modularity through  
533 symbolic and object-centric reasoning while supporting human-in-the-loop control at the skill level,  
534 offering a promising path toward transparent, adaptable, and aligned agents.

535  
536 **Reproducibility Statement.** We have taken several measures to ensure reproducibility of our  
537 results. Details of the proposed method, including model architectures, training procedures, evaluation  
538 protocols, hyperparameters, implementation details and LLM interactions are provided in the main  
539 paper and the appendix. Additionally, we release source code and configuration files to reproduce all  
experiments, along with environment setup instructions.

540 REFERENCES  
541

542 Kamal Acharya, Waleed Raza, Carlos M. J. M. Dourado Júnior, Alvaro Velasquez, and Houbing Her-  
543 bert Song. Neurosymbolic Reinforcement Learning and Planning: A Survey. *IEEE Transactions*  
544 *on Artificial Intelligence*, 5(5):1939–1953, 2024.

545 Riad Akrou, Davide Tateo, and Jan Peters. Towards reinforcement learning of human readable  
546 policies. In *The European Conference on Machine Learning and Principles and Practice of*  
547 *Knowledge Discovery in Databases: The 1st Workshop on Deep Continuous-Discrete Machine*  
548 *Learning*, 2019.

549 Greg Anderson, Abhinav Verma, İşıl Dillig, and Swarat Chaudhuri. Neurosymbolic Reinforcement  
550 Learning with Formally Verified Exploration. In *Neural Information Processing Systems*, volume  
551 abs/2009.12612, 2020.

553 Jacob Andreas, Dan Klein, and Sergey Levine. Modular Multitask Reinforcement Learning with  
554 Policy Sketches. In *International Conference on Machine Learning (ICML)*, pp. 166–175, 2017.

556 Pierre-Luc Bacon, Jean Harb, and Doina Precup. The Option-Critic Architecture. In *AAAI Conference*  
557 *on Artificial Intelligence (AAAI)*, pp. 1726–1734, 2017.

558 Adrià Puigdomènech Badia, Bilal Piot, Steven Kapturowski, Pablo Sprechmann, Alex Vitvitskyi,  
559 Zhaohan Daniel Guo, and Charles Blundell. Agent57: Outperforming the Atari Human Benchmark.  
560 In *International Conference on Machine Learning (ICML)*, pp. 507–517, 2020.

562 Andrew G Barto and Sridhar Mahadevan. Recent advances in hierarchical reinforcement learning.  
563 *Discrete event dynamic systems*, 13:341–379, 2003.

564 Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The Arcade Learning  
565 Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence*  
566 *Research (JAIR)*, 47:253–279, 2013.

568 Brett Daley and Christopher Amato. Reconciling  $\lambda$ -Returns with Experience Replay. In *Conference*  
569 *on Neural Information Processing Systems (NeurIPS)*, 2019.

570 Quentin Delfosse, Hikaru Shindo, Devendra Singh Dhami, and Kristian Kersting. Interpretable and  
571 Explainable Logical Policies via Neurally Guided Symbolic Abstraction. In *Conference on Neural*  
572 *Information Processing Systems (NeurIPS)*, 2023a.

574 Quentin Delfosse, Wolfgang Stammer, Thomas Rothenbächer, Dwarak Vittal, and Kristian Kersting.  
575 *Boosting Object Representation Learning via Motion and Object Continuity*. Springer Nature  
576 Switzerland, 2023b.

577 Quentin Delfosse, Jannis Blüml, Bjarne Gregori, Sebastian Sztwiertnia, and Kristian Kersting. Ocatari:  
578 Object-Centric Atari 2600 Reinforcement Learning Environments. *Reinforcement Learning*  
579 *Conference (RLC)*, 1:400–449, 2024a.

581 Quentin Delfosse, Sebastian Sztwiertnia, Mark Rothermel, Wolfgang Stammer, and Kristian Kersting.  
582 Interpretable Concept Bottlenecks to Align Reinforcement Learning Agents. In *Conference on*  
583 *Neural Information Processing Systems (NeurIPS)*, volume abs/2401.05821, 2024b.

585 Quentin Delfosse, Jannis Blüml, Fabian Tatai, Théo Vincent, Bjarne Gregori, Elisabeth Dillies, Jan  
586 Peters, Constantin Rothkopf, and Kristian Kersting. Deep reinforcement learning agents are not  
587 even close to human intelligence. *arXiv preprint*, 2025.

588 Thomas G. Dietterich. Hierarchical Reinforcement Learning with the MAXQ Value Function  
589 Decomposition. *Journal of Artificial Intelligence Research (JAIR)*, 13:227–303, 2000.

590 Elisabeth Dillies, Quentin Delfosse, Jannis Blüml, Raban Emunds, Florian Peter Busch, and Kristian  
591 Kersting. Better decisions through the right causal world model. 2025.

593 Matteo Gallici, Mattie Fellows, Benjamin Ellis, B. Pou, Ivan Masmitja, J. Foerster, and Mario Martin.  
594 Simplifying Deep Temporal Difference Learning. *arXiv.org*, abs/2407.04811, 2024.

594 Claire Glanois, Paul Weng, Matthieu Zimmer, Dong Li, Tianpei Yang, Jianye Hao, and Wulong Liu.  
 595 A survey on interpretable reinforcement learning. *Machine Learning*, 113(8):5847–5890, 2024.  
 596

597 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy  
 598 maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference  
 599 on Machine Learning (ICML)*. Pmlr, 2018.

600 Danijar Hafner. Benchmarking the spectrum of agent capabilities. *arXiv preprint arXiv:2109.06780*,  
 601 2021.  
 602

603 Rishi Hazra and Luc De Raedt. Deep Explainable Relational Reinforcement Learning: A Neuro-  
 604 Symbolic Approach. In *European Conference on Machine Learning and Knowledge Discovery in  
 605 Databases (PKDD)*, pp. 213–229, 2023.

606 Shengyi Huang, Rousslan Fernand Julien Dossa, Chang Ye, Jeff Braga, Dipam Chakraborty, Kinal  
 607 Mehta, and João G.M. Araújo. Cleanrl: High-quality single-file implementations of deep rein-  
 608 force learning algorithms. *Journal of Machine Learning Research*, 23(274):1–18, 2022. URL  
 609 <http://jmlr.org/papers/v23/21-1342.html>.  
 610

611 Steven James, Benjamin Rosman, and George Konidaris. Autonomous Learning of Object-Centric  
 612 Abstractions for High-Level Planning. In *International Conference on Learning Representations  
 613 (ICLR)*, 2022.

614 Mu Jin, Zhihao Ma, Kebing Jin, Hankz Hankui Zhuo, Chen Chen, and Chao Yu. Creativity of AI:  
 615 Automatic Symbolic Option Discovery for Facilitating Deep Reinforcement Learning. In *AAAI  
 616 Conference on Artificial Intelligence (AAAI)*, pp. 7042–7050, 2022.  
 617

618 Daniel Kahneman. *Thinking, fast and slow*. macmillan, 2011.  
 619

620 Timo Kaufmann, Jannis Blüml, Antonia Wüst, Quentin Delfosse, Kristian Kersting, and Eyke Hüller-  
 621 meier. Ocalm: Object-centric assessment with language models. *arXiv preprint arXiv:2406.16748*,  
 622 2024.

623 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloé Rolland, Laura Gustafson, Tete  
 624 Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross B. Girshick.  
 625 Segment Anything. In *IEEE International Conference on Computer Vision (ICCV)*, pp. 3992–4003,  
 626 2023.  
 627

628 Hector Kohler, Quentin Delfosse, R. Akour, Kristian Kersting, and Philippe Preux. Interpretable  
 629 and Editable Programmatic Tree Policies for Reinforcement Learning. *arXiv.org*, abs/2405.14956,  
 630 2024.

631 Matteo Leonetti, Luca Iocchi, and Peter Stone. A synthesis of automated planning and reinforcement  
 632 learning for efficient, robust decision-making. *Artificial Intelligence*, 241:103–130, 2016.  
 633

634 Yuezhang Li, Katia P. Sycara, and Rahul Iyer. Object-sensitive Deep Reinforcement Learning. In  
 635 *Global Conference on Artificial Intelligence (GCAI)*, pp. 20–35, 2017.  
 636

637 Amarildo Likmeta, Alberto Maria Metelli, Andrea Tirinzoni, Riccardo Giol, Marcello Restelli, and  
 638 Danilo Romano. Combining reinforcement learning with rule-based controllers for transparent and  
 639 general decision-making in autonomous driving. *Robotics and Autonomous Systems*, 131:103568,  
 640 2020.

641 Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa,  
 642 David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In  
 643 *International Conference on Learning Representations (ICLR)*, 2016.  
 644

645 Zhixuan Lin, Yi-Fu Wu, Skand Vishwanath Peri, Weihao Sun, Gautam Singh, Fei Deng, Jindong  
 646 Jiang, and Sungjin Ahn. Space: Unsupervised Object-Oriented Scene Representation via Spatial  
 647 Attention and Decomposition. In *International Conference on Learning Representations (ICLR)*,  
 2020.

648 Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold,  
 649 Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-Centric Learning with Slot  
 650 Attention. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.

651

652 Chris Lu, Jakub Kuba, Alistair Letcher, Luke Metz, Christian Schroeder de Witt, and Jakob Foerster.  
 653 Discovered policy optimisation. *Advances in Neural Information Processing Systems*, 35:16455–  
 654 16468, 2022.

655 Lirui Luo, Guoxi Zhang, Hongming Xu, Yaodong Yang, Cong Fang, and Qing Li. End-to-End  
 656 Neuro-Symbolic Reinforcement Learning with Textual Explanations. In *International Conference  
 657 on Machine Learning (ICML)*, 2024.

658

659 Daoming Lyu, Fangkai Yang, Bo Liu, and Steven Gustafson. Sdrl: Interpretable and Data-Efficient  
 660 Deep Reinforcement Learning Leveraging Symbolic Planning. In *AAAI Conference on Artificial  
 661 Intelligence (AAAI)*, pp. 2970–2977, 2019.

662

663 Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman,  
 664 Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding  
 665 large language models. In *The Twelfth International Conference on Learning Representations*,  
 666 2024. URL <https://openreview.net/forum?id=IEduRU055F>.

667

668 Marlos C. Machado, Marc G. Bellemare, Erik Talvitie, Joel Veness, Matthew J. Hausknecht, and  
 669 Michael Bowling. Revisiting the arcade learning environment: Evaluation protocols and open  
 670 problems for general agents (extended abstract). In *Proceedings of the Twenty-Seventh International  
 Joint Conference on Artificial Intelligence*, 2018.

671

672 Francis Maes, Raphael Fonteneau, Louis Wehenkel, and Damien Ernst. Policy Search in a Space of  
 673 Simple Closed-form Formulas: Towards Interpretability of Reinforcement Learning. In *International  
 Conference on Discovery Science (DS)*, pp. 37–51, 2012.

674

675 Michael Matthews, Michael Beukman, Benjamin Ellis, Mikayel Samvelyan, Matthew Jackson,  
 676 Samuel Coward, and Jakob Foerster. Craftax: A Lightning-Fast Benchmark for Open-Ended Rein-  
 677 force Learning. In *International Conference on Machine Learning*, volume abs/2402.16801,  
 678 2024.

679

680 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bel-  
 681 lemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen,  
 682 Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra,  
 683 Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning.  
 684 *Nature*, 518(7540):529–533, 2015.

685

686 Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham  
 687 Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev  
 688 Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollar, and Christoph Feichtenhofer.  
 689 Sam 2: Segment Anything in Images and Videos. In *The Thirteenth International Conference on  
 Learning Representations*, volume abs/2408.00714, 2025.

690

691 Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use  
 692 interpretable models instead. *Nature machine intelligence*, 1(5):206–215, 2019.

693

694 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy  
 695 Optimization Algorithms. *arXiv.org*, abs/1707.06347, 2017.

696

697 Mohit Sharma, Jacky Liang, Jialiang Zhao, Alex LaGrassa, and Oliver Kroemer. Learning to  
 698 Compose Hierarchical Object-Centric Controllers for Robotic Manipulation. In *Conference on  
 699 Robot Learning (CoRL)*, pp. 822–844, 2020.

700

701 Ibne Farabi Shihab, Sanjeda Akter, and Anuj Sharma. Detecting and mitigating reward hacking in  
 702 reinforcement learning systems: A comprehensive empirical study. *arXiv preprint*, 2025.

703

704 Hikaru Shindo, Quentin Delfosse, Devendra Singh Dhami, and Kristian Kersting. Blendrl: A  
 705 framework for merging symbolic and neural policy learning. In *The Thirteenth International  
 Conference on Learning Representations*, 2025.

702 Andrew Silva, Taylor W. Killian, I. D. Rodriguez, Sung-Hyun Son, and M. Gombolay. Optimization  
 703 Methods for Interpretable Differentiable Decision Trees in Reinforcement Learning. *arXiv: Learning*, 2019.  
 704

705 Dmitriy Smirnov, Michaël Gharbi, Matthew Fisher, Vitor Guizilini, Alexei A. Efros, and Justin M.  
 706 Solomon. Marionette: Self-Supervised Sprite Learning. In *Conference on Neural Information*  
 707 *Processing Systems (NeurIPS)*, pp. 5494–5505, 2021.  
 708

709 Richard S. Sutton, Doina Precup, and Satinder Singh. Between MDPs and Semi-MDPs: A Framework  
 710 for Temporal Abstraction in Reinforcement Learning. *Artificial Intelligence*, 112(1-2):181–211,  
 711 1999.  
 712

713 Abhinav Verma, Hoang Minh Le, Yisong Yue, and Swarat Chaudhuri. Imitation-Projected Pro-  
 714 grammatic Reinforcement Learning. In *Conference on Neural Information Processing Systems*  
 715 (*NeurIPS*), pp. 15726–15737, 2019.  
 716

717 Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David  
 718 Silver, and Koray Kavukcuoglu. Feudal Networks for Hierarchical Reinforcement Learning. In *International Conference on Machine Learning (ICML)*, pp. 3540–3549, 2017.  
 719

720 Yue Wu, Yewen Fan, Paul Pu Liang, Amos Azaria, Yuanzhi Li, and Tom M. Mitchell. Read and  
 721 Reap the Rewards: Learning to Play Atari with the Help of Instruction Manuals. In *Conference on*  
 722 *Neural Information Processing Systems (NeurIPS)*, 2023.  
 723

724 Tianbao Xie, Siheng Zhao, Chen Henry Wu, Yitao Liu, Qian Luo, Victor Zhong, Yanchao Yang, and  
 725 Tao Yu. Text2reward: Reward shaping with language models for reinforcement learning. In *The*  
 726 *Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May*  
 727 *7-11, 2024, 2024*.  
 728

729 Fangkai Yang, Daoming Lyu, Bo Liu, and Steven Gustafson. Peorl: Integrating Symbolic Planning  
 730 and Hierarchical Reinforcement Learning for Robust Decision-Making. In *International Joint*  
 731 *Conference on Artificial Intelligence (IJCAI)*, pp. 4860–4866, 2018.  
 732

733 Yupei Yang, Biwei Huang, Fan Feng, Xinyue Wang, Shikui Tu, and Lei Xu. Towards generalizable  
 734 reinforcement learning via causality-guided self-adaptive representations. In *The Thirteenth*  
 735 *International Conference on Learning Representations*, 2025.  
 736

737 Zihan Ye, O. Arenz, and Kristian Kersting. Learning from Less: Guiding Deep Reinforcement  
 738 Learning with Differentiable Symbolic Planning. *arXiv*, 2025.  
 739

740 Tom Zahavy, Nir Ben-Zrihem, and Shie Mannor. Graying the black box: Understanding DQNs. In  
 741 *International Conference on Machine Learning (ICML)*, pp. 1899–1908, 2016.  
 742

743 Chiyuan Zhang, O. Vinyals, R. Munos, and Samy Bengio. A Study on Overfitting in Deep Reinforce-  
 744 ment Learning. *arXiv.org*, abs/1804.06893, 2018.  
 745

## A LLM USAGE

747 Beyond integration within the pipeline (cf. Section G), LLMs were used solely to improve text  
 748 readability and generate boilerplate code.  
 749

## B MOTIVATION - INTERPRETABLE POLICY SIZES

750 We visualize the actual logic policies retrieved from existing interpretable RL methods in Listing 1  
 751 and Listing 2. While these policies have maximum transparency, they are difficult to interpret for  
 752 humans due to the massive size. NEXUS on the other hand presents a simple solution by incorporating  
 753 hierarchical abstraction and thus allowing for compact and truly interpretable policies, cf. Figure 5.  
 754

```

756
757 def play( state ):
758     if state .Ball_0.prev_x <= 1.24:
759         if state .Ball_0.y - state .Enemy_0.prev_y <= -0.41:
760             if state .Ball_0.x - state .Ball_0.prev_x <= 0.09:
761                 if state .Player_0.y - state .Enemy_0.y <= -0.91:
762                     if state .Ball_0.x <= 0.73:
763                         if state .Player_0.y - state .Ball_0.x <= -1.20:
764                             return "LEFT"
765                         else:
766                             return "RIGHT"
767                     else:
768                         return "NOOP"
769             else:
770                 if state .Ball_0.x <= -0.05:
771                     return "NOOP"
772             else:
773                 if state .Player_0.y - state .Player_0.prev_y <= -0.19:
774                     return "NOOP"
775                 else:
776                     if state .Ball_0.prev_x - state .Ball_0.prev_y <= -0.38:
777                         return "NOOP"
778                     else:
779                         if state .Ball_0.x - state .Ball_0.prev_x <= -0.09:
780                             return "RIGHT"
781                     else:
782                         return "RIGHT"
783             else:
784                 if state .Ball_0.x - state .Enemy_0.y <= -1.46:
785                     if state .Player_0.y - state .Ball_0.prev_x <= -0.92:
786                         return "LEFT"
787                     else:
788                         return "RIGHT"
789             else:
790                 if state .Player_0.y - state .Ball_0.y <= -0.63:
791                     if state .Ball_0.prev_x <= -0.05:
792                         return "LEFT"
793                     else:
794                         if state .Player_0.y - state .Ball_0.y <= -0.80:
795                             return "LEFT"
796                         else:
797                             if state .Ball_0.prev_x <= 0.80:
798                                 return "LEFT"
799                             else:
800                                 return "NOOP"
801             else:
802                 if state .Ball_0.x <= -0.07:
803                     if state .Ball_0.prev_y <= -0.15:
804                         if state .Ball_0.prev_y <= -1.96:
805                             return "LEFT"
806                         else:
807                             return "LEFT"
808             else:
809                 if state .Ball_0.y <= 1.24:
810                     if state .Player_0.y - state .Ball_0.y <= -0.48:

```

```

810
811         if state.Ball_0.prev_x <= 0.68:
812             return "NOOP"
813         else:
814             return "NOOP"
815         else:
816             if state.Ball_0.x <= 0.54:
817                 if state.Player_0.y <= -0.82:
818                     return "NOOP"
819                 else:
820                     if state.Enemy_0.y - state.Enemy_0.prev_y <=
821                         -0.09:
822                         return "RIGHT"
823                     else:
824                         return "RIGHT"
825             else:
826                 if state.Ball_0.x - state.Ball_0.y <= -0.21:
827                     return "NOOP"
828                 else:
829                     return "RIGHT"
830             else:
831                 if state.Ball_0.y <= 1.54:
832                     return "NOOP"
833             else:
834                 if state.Player_0.y - state.Ball_0.y <= -0.54:
835                     if state.Player_0.y - state.Ball_0.y <= -0.88:
836                         return "LEFT"
837                     else:
838                         if state.Ball_0.y <= -1.40:
839                             return "RIGHT"
840                         else:
841                             return "NOOP"
842             else:
843                 if state.Enemy_0.y - state.Enemy_0.prev_y <= 0.93:
844                     return "RIGHT"
845                 else:
846                     return "LEFT"

```

Listing 1: Pong policy of SCoBots (Delfosse et al., 2024b)

```

847 up_air(X):-oxygen_low(B).
848 up_divers_collected (X):- all_divers_collected (D).
849 fire_left (X):-same_depth_enemy(P,E),visible_enemy(E),facing_left (P),right_of_enemy(P,E).
850 fire_right (X):-same_depth_enemy(P,E),visible_enemy(E),facing_right (P),left_of_enemy(P,E)
851 .
852 left_aim (X):-right_of_enemy(P,E), facing_right (P), same_depth_enemy(P,E),visible_enemy(E).
853 right_aim (X):-left_of_enemy(P,E), facing_left (P), same_depth_enemy(P,E),visible_enemy(E).
854 down_aim(X):-higher_than_enemy(P,E),visible_enemy(E).
855 up_aim(X):-deeper_than_enemy(P,E),visible_enemy(E).
856 up_evade(X):-close_by_enemy(P,E),same_depth_enemy(P,E),visible_enemy(E).
857 down_evade(X):-close_by_enemy(P,E),same_depth_enemy(P,E),visible_enemy(E).
858 up_evade(X):-close_by_missile (P,M),same_depth_missile(P,M), visible_missile (M).
859 down_evade(X):-close_by_missile(P,M),same_depth_missile(P,M), visible_missile (M).
860 left_to_diver (X):- right_of_diver (P,D),close_by_diver (P,D), visible_diver (D).
861 right_to_diver (X):- left_of_diver (P,D),close_by_diver (P,D), visible_diver (D).
862 up_to_diver (X):-deeper_than_diver (P,D),close_by_diver (P,D), visible_diver (D).
863 down_to_diver(X):-higher_than_diver (P,D),close_by_diver (P,D), visible_diver (D).

```

Listing 2: Seaquest policy of NUDGE (Delfosse et al., 2023a)

```

864
865 logits_noop1 = -0.56*y_agent_1**2 - 0.38*y_agent_1*y_agent_2 - 0.087*y_agent_1*
866     y_opponent_1 - 0.16*y_agent_1*y_opponent_2 - 0.76*y_agent_1*y_opponent_3 - 0.51*
867     y_agent_1*y_opponent_4 - 0.54*y_agent_1 - 0.24*y_agent_2**2 - 0.073*y_agent_2 +
868     0.27*y_agent_4**2 + 0.55*y_agent_4 - 0.078*y_opponent_1**2 - 0.33*y_opponent_1*
869     y_opponent_2 - 0.2*y_opponent_1 - 0.35*y_opponent_2**2 - 0.5*y_opponent_2 - 0.34*
870     y_opponent_3**2 - 0.45*y_opponent_3*y_opponent_4 - 0.32*y_opponent_3 - 0.15*
871     y_opponent_4**2 - 0.19*y_opponent_4 + 1.1
872 logits_noop2 = -0.074*y_agent_1*y_opponent_2 + 0.059*y_agent_1*y_opponent_3 - 0.097*
873     y_agent_4 - 0.16*y_opponent_1*y_opponent_2 - 0.18*y_opponent_2**2 - 0.27*
874     y_opponent_2 + 0.063*y_opponent_4
875 logits_up1 = 0.23*y_agent_1**2 + 0.59*y_agent_1*y_agent_2 + 0.4*y_agent_2**2 + 0.11*
876     y_agent_2 - 1.5*y_agent_4**2 - 3.6*y_agent_4 + 0.068*y_opponent_3 + 1.1
877 logits_down1 = 0.09*x_ball_3 + 0.12*x_ball_4 - 0.21*y_agent_1**2 + 0.12*y_agent_1*
878     y_opponent_1 + 0.27*y_agent_1*y_opponent_2 - 0.43*y_agent_1*y_opponent_3 - 0.28*
879     y_agent_1*y_opponent_4 + 0.13*y_agent_2 + 0.14*y_agent_4**2 + 0.43*y_agent_4 +
880     0.087*y_ball_3 + 0.15*y_ball_4 + 0.14*y_opponent_1**2 + 0.6*y_opponent_1*
881     y_opponent_2 + 0.61*y_opponent_1 + 0.65*y_opponent_2**2 + 1.1*y_opponent_2 - 0.2*
882     y_opponent_3**2 - 0.26*y_opponent_3*y_opponent_4 - 2.8*y_opponent_3 - 0.085*
883     y_opponent_4**2 - 0.14*y_opponent_4 - 2.3
884 logits_up2 = 0.063*x_ball_4 - 0.078*y_agent_1 + 0.18*y_agent_2**2 + 0.52*y_agent_2*
885     y_agent_3 + 0.35*y_agent_2*y_opponent_1 + 0.29*y_agent_2*y_opponent_2 + 0.26*
886     y_agent_2 + 0.38*y_agent_3**2 + 0.51*y_agent_3*y_opponent_1 + 0.42*y_agent_3*
887     y_opponent_2 + 1.6*y_agent_3 - 8.2*y_agent_4 - 0.085*y_ball_3 + 0.17*y_opponent_1**2
888     + 0.28*y_opponent_1*y_opponent_2 + 0.25*y_opponent_1 + 0.11*y_opponent_2**2 +
889     0.15*y_opponent_2 - 0.074*y_opponent_3 + 0.26 logits_down2 = -0.052*x_ball_1 - 0.068*
890     x_ball_3 - 0.093*x_ball_4 + 0.18*y_agent_1 - 0.17*y_agent_2**2 - 0.49*y_agent_2*
891     y_agent_3 - 0.33*y_agent_2*y_opponent_1 - 0.27*y_agent_2*y_opponent_2 - 0.39*
892     y_agent_2 - 0.35*y_agent_3**2 - 0.48*y_agent_3*y_opponent_1 - 0.4*y_agent_3*
893     y_opponent_2 - 0.38*y_agent_3 + 0.15*y_agent_4**2 + 0.54*y_agent_4 - 0.06*y_ball_1 -
894     0.064*y_ball_3 - 0.11*y_ball_4 - 0.17*y_opponent_1**2 - 0.28*y_opponent_1*
895     y_opponent_2 - 0.58*y_opponent_1 - 0.13*y_opponent_2**2 - 0.38*y_opponent_2 + 2.2*
896     y_opponent_3 - 0.052*y_opponent_4 - 3.6
897
898 action_noop = [exp(logits_noop1) + exp(logits_noop2)] / sum(exp(logits))
899 action_up = [exp(logits_up1) + exp(logits_up2)] / sum(exp(logits))
900 action_down = [exp(logits_down1) + exp(logits_down2)] / sum(exp(logits))

```

Listing 3: Pong policy of INSIGHT (Luo et al., 2024)

## C HIERARCHICAL PQN ALGORITHM

The complete algorithm for hierarchical PQN with a neural meta-policy is provided in Algorithm 2.

## D IMPLEMENTATION DETAILS

**Setup.** We base our implementations on CleanRL (Huang et al., 2022) and PureJaxRL (Lu et al., 2022), adapting them to object-centric inputs by replacing convolutional encoders with lightweight MLPs for feature extraction. Hyperparameters are listed in Table 1 and Table 2 and remain largely consistent with the original implementations, except for an increased number of parallel environments enabled by the efficiency of JAX-based code.

The exploration parameter  $\epsilon$  was selected via a brief hyperparameter sweep in the range  $[1, 0.001]$ , using final test return as the selection criterion.

Each experiment is run with three random seeds (0, 1, 2) to ensure reproducibility. Reported plots include the corresponding standard deviation.

All experiments were conducted on a single NVIDIA Tesla V100-SXM3-32GB-H GPU on an NVIDIA DGX Server (Version 5.1.0) with CUDA 12.4.

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928

---

**Algorithm 2** Hierarchical PQN

---

929  
930 **Require:** Update period  $U$ , number of parallel environments  $E$ , number of skills  $N$ , exploration  
931 probability  $\epsilon$   
932 **Ensure:** Learned Q-network parameters  $\{\phi_n\}_{n=1}^N, \phi_{\text{meta}}$   
933 1: Initialize Q-network parameters  $\{\phi_n\}_{n=1}^N, \phi_{\text{meta}}$   
934 2: Sample initial states  $s_0^e \sim P_0$  for  $e \in \{0, \dots, E-1\}$   
935 3:  $t \leftarrow 0$   
936 4: **for** each episode **do**  
937 5:   **for all**  $e \in \{0, \dots, E-1\}$  **in parallel do**  
938 6:     Sample skill  $\pi_t^e \sim \pi_{\text{meta}}$   
939 7:     With probability  $\epsilon$ :  $a_t^e \sim \text{Unif}$ , else  $a_t^e \sim \pi_t^e$   
940 8:     Sample rewards  $r_t^e \sim P_R(s_t^e, a_t^e)$ , skill rewards  $r_t^{e,n} \sim P_{R,n}(s_t^e, a_t^e)$  for all  $n$   
941 9:     Sample next state  $s_{t+1}^e \sim P_S(s_t^e, a_t^e)$   
942 10:     $t \leftarrow t + 1$   
943 11: **end for**  
944 12: **if**  $t \bmod U = 0$  **then**  
945 13:    Compute meta  $\lambda$ -returns  $R_{\lambda,t-1}^e$  to  $R_{\lambda,t-U}^e$  for all  $e$   
946 14:    Compute skill  $\lambda$ -returns  $R_{\lambda,t-1}^{e,n}$  to  $R_{\lambda,t-U}^{e,n}$  for all  $e, n$   
947 15:    **for** number of epochs **do**  
948 16:     **for** number of minibatches **do**  
949 17:       Sample minibatch  $B$  of size  $b \leq EU$  from  $\{(t-U, 0), \dots, (t-1, E-1)\}$   
950 18:       Update meta:  
951       
$$\phi_{\text{meta}} \leftarrow \phi_{\text{meta}} + \frac{\alpha_t}{2b} \nabla_{\phi_{\text{meta}}} \sum_{(j,\tau) \in B} \left( R_{\lambda,\tau}^j - Q_{\phi_{\text{meta}}}(s_{\tau}^j) \right)^2$$
  
952  
953 19:     Update skills:  
954     
$$\phi_n \leftarrow \phi_n + \frac{\alpha_t}{2b} \nabla_{\phi_n} \sum_{(j,\tau) \in B} \left( R_{\lambda,\tau}^{j,n} - Q_{\phi_n}(s_{\tau}^{j,n}) \right)^2, \quad \forall n \in \{1, \dots, N\}$$
  
955  
956  
957 20:    **end for**  
958 21:    **end for**  
959 22:   **end if**  
960 23: **end for**

---

962  
963  
964  
965  
966  
967  
968  
969  
970  
971



Figure 10: The evaluation environments: Crafter, Seaquest and Kangaroo.

We provide an anonymized code repository<sup>2</sup> that includes all necessary code and config files to reproduce the results from the experiments section, including plots.

## D.1 ENVIRONMENTS

Screenshots of the environments Crafter, Seaquest and Kangaroo are available in Figure 10.

## D.2 EVALUATION METRICS

We empirically evaluate agent performance using three metrics: (1) **Skill Returns** to test whether the skills were learned successfully, (2) **Human-Normalized Score** (HNS) for absolute performance relative to human and random baselines and (3) **Aligned Environment Goal Scores** that measure performance based on the main goals described in the game’s manual.

**Human Normalized Score.** Human-Normalized Score standardizes agent performance across Atari environments by accounting for differences in reward scales (Mnih et al., 2015; Machado et al., 2018). Given the average agent score  $A$ , human score  $H$ , and random score  $R$ , HNS is defined as:

$$\text{HNS} = \frac{A - R}{|H - R|}$$

A value of 1.0 indicates human-level performance, values greater than 1.0 indicate superhuman performance, and values below 0 denote sub-random behavior. We adopt the human and random baselines from Badia et al. (2020), derived from professional human play.

**Aligned Environment Goal Scores.** Environment reward signals may not always align with the intended task objectives and can be susceptible to reward hacking. In such cases, agents may learn high-reward behaviors that are non-intuitive and deviate from human-like solutions. To better capture progress toward the actual environment goals, we define two aligned goal-based metrics grounded in the objectives stated in the game manuals.

For *Seaquest*, the goal is to rescue as many divers as possible; for *Kangaroo*, it is to help the mother kangaroo reach and rescue her baby, located on the topmost platform. Accordingly, we track the number of divers retrieved and the number of platforms reached, respectively.

## E ADDITIONAL RESULTS

To assess the generality of NEXUS, we extend our evaluation to three additional Atari environments: *Pong*, *Breakout*, and *Freeway*. Skill learning curves for these environments are shown in Figure 11,

<sup>2</sup>[https://anonymous.4open.science/r/symbolic\\_options-302C/](https://anonymous.4open.science/r/symbolic_options-302C/)

Table 1: Atari Hyperparameters

Parameter	PQN(-based)	PPO
Total Timesteps	$5 \times 10^7$	$5 \times 10^7$
Num Environments	1024	128
Num Steps per Update	128	128
Learning Rate	$1.0 \times 10^{-4}$	$2.5 \times 10^{-4}$
Max Grad Norm	10	0.5
Discount Factor ( $\gamma$ )	0.99	0.99
GAE Lambda ( $\lambda$ )	0.65 (0.5 with learned meta)	0.95
GAE Meta Lambda ( $\lambda$ )	0.9	–
Num Epochs	5	2
Num Minibatches	128	4
Hidden Size	64	–
Num Layers	3	–
Normalization	Layer Norm	–
Clip $\epsilon$	–	0.2
Entropy Coef	–	0.01
Value Function Coef	–	0.5
Anneal LR	False	True
$\epsilon$ -Start/End/Decay	1.0 / 0.1 / 0.3	–
Meta $\epsilon$ -Start/End/Decay	1.0 / 0.001 / 0.3	–

Table 2: Crafter Hyperparameters

Parameter	PQN(-based)
Total Timesteps	$1 \times 10^9$
Num Environments	512
Num Steps per Update	128
Learning Rate	$3.0 \times 10^{-4}$
Max Grad Norm	0.5
Discount Factor ( $\gamma$ )	0.99
GAE Lambda ( $\lambda$ )	0.5
GAE Meta Lambda ( $\lambda$ )	0.5
Num Epochs	4
Num Minibatches	4
Hidden Size	512
Num Layers	1
Normalization	Layer Norm
Anneal LR	True
$\epsilon$ -Start/End/Decay	1.0 / 0.005 / 0.1
Meta $\epsilon$ -Start/End/Decay	1.0 / 0.005 / 0.1

while comparisons to baseline agents and ablations—evaluated via human-normalized scores—are presented in Figure 12. Note that these games are generally less complex than *Kangaroo* and *Seaquest*, and the learned skills are not strictly necessary to achieve the environment goals. In particular, the skills "Move Up" and "Avoid Crash" in *Freeway* largely correspond to atomic actions such as `forward` or `noop`. As such, a simple fixed meta-policy operating directly on primitive actions could suffice for solving this task.

**Neuro-symbolic RL baselines.** For improved comparison, we also provide baseline scores of existing interpretable and neuro-symbolic methods on both the default games (cf. Figure 13a) and the simplifications (cf. Figure 13b). The methods are NUDGE (Delfosse et al., 2023a) and BlendRL (Shindo et al., 2025). We also provide a tabular overview of our results and also add the reported scores from SCoBots (Delfosse et al., 2024b) (game simplification scores from Delfosse et al. (2025) in Table 3).

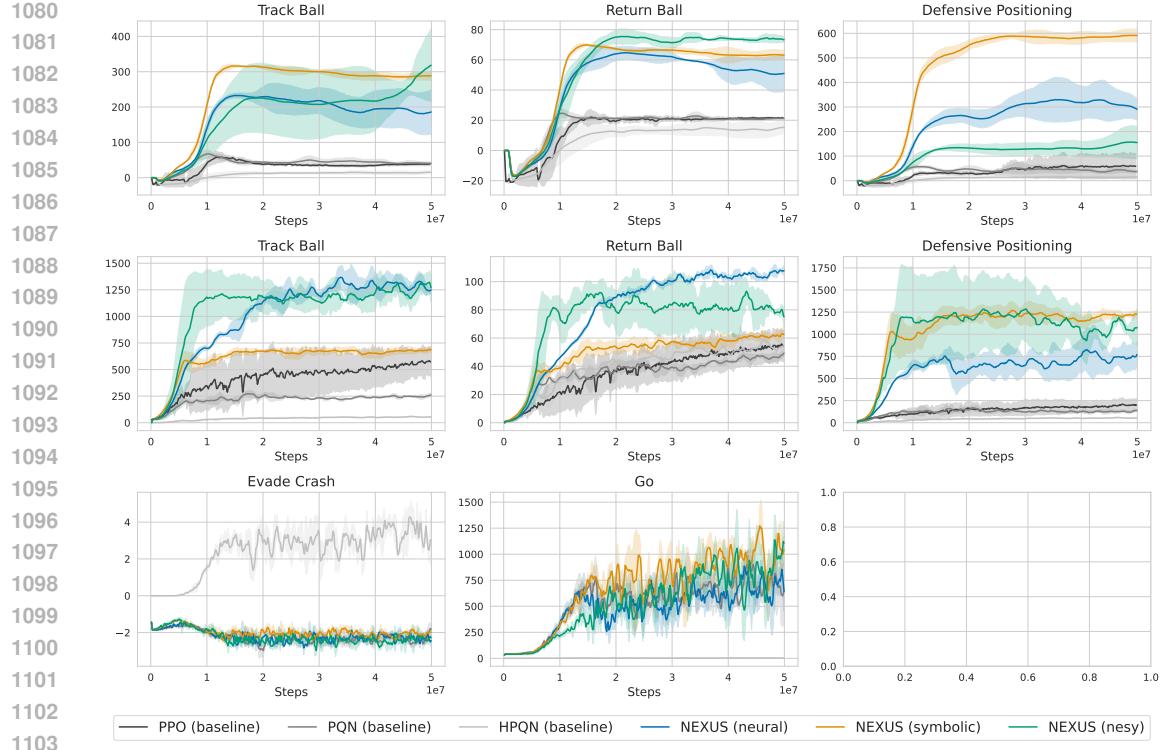
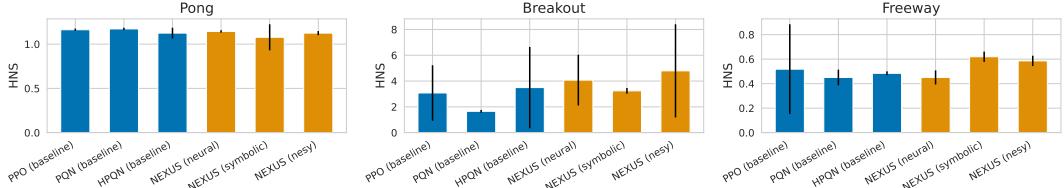
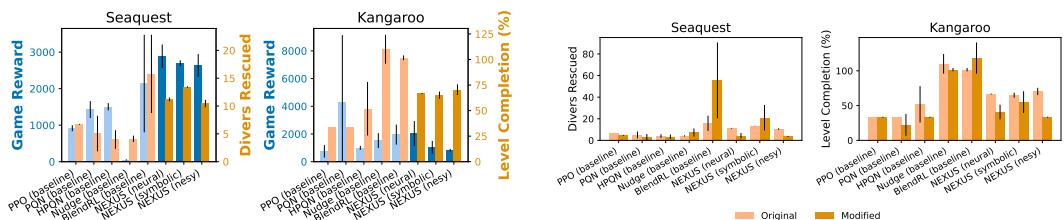
Figure 11: NEXUS successfully learns skills in *Pong* (top), *Breakout* (middle), and *Freeway* (bottom).

Figure 12: NEXUS achieves performance comparable to baseline methods across the three evaluated environments.

## F NOISY DETECTIONS

We evaluate the robustness of NEXUS on noise in the object detections both during training and testing. We conduct experiments with 5% and 10% misdetection rate. To emulate the effect of a kalman filter estimating the objects movements during each step, instead of zeroing out the detected objects, we keep the previous time step detections. On top of that, we add gaussian noise with a standard deviation of 3 at each attribute. We show the results with the noisy detection during training in Figure 14. Experiencing noise only during testing is evaluated in Figure 15.



(a) Comparison to other interpretable RL methods on default environments.

(b) Comparison to other interpretable RL methods on simplified games.

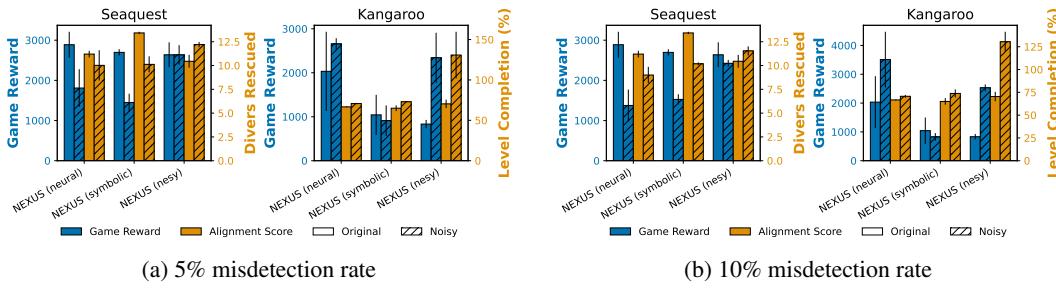
Figure 13: Additional neuro-symbolic RL baselines.

1134  
1135  
1136  
1137  
1138  
1139

Table 3: Performance Comparison on Seaquest and Kangaroo

Algorithm	Seaquest				Kangaroo			
	Default	Divers Collected	Simplification	Divers Collected	Default	Level Completion (%)	Simplification	Level Completion (%)
PPO	915.3 $\pm$ 82.5	(6.7 $\pm$ 0.1)	0.0 $\pm$ 0.0	(4.7 $\pm$ 0.5)	750.0 $\pm$ 450.0	(33.3 $\pm$ 0.0)	0.0 $\pm$ 0.0	(33.3 $\pm$ 0.0)
PQN	1428.3 $\pm$ 232.1	(5.1 $\pm$ 3.2)	0.0 $\pm$ 0.0	(3.0 $\pm$ 2.8)	<b>4300.0 <math>\pm</math> 4833.2</b>	(33.3 $\pm$ 0.0)	200.0 $\pm$ 141.4	(22.2 $\pm$ 15.7)
NEXUS (neural)	<b>2887.6 <math>\pm</math> 322.4</b>	(11.2 $\pm$ 0.3)	0.0 $\pm$ 0.0	(4.3 $\pm$ 2.5)	2033.3 $\pm$ 899.4	(66.7 $\pm$ 0.0)	666.7 $\pm$ 94.3	(40.7 $\pm$ 10.5)
NEXUS (symbolic)	2697.1 $\pm$ 77.7	(13.4 $\pm$ 0.1)	1853.3 $\pm$ 2621.0	(21.0 $\pm$ 11.8)	1042.9 $\pm$ 454.9	(65.0 $\pm$ 3.6)	1346.5 $\pm$ 1042.9	(55.0 $\pm$ 15.6)
NEXUS (nesy)	2637.1 $\pm$ 309.9	(10.4 $\pm$ 0.7)	0.0 $\pm$ 0.0	(3.7 $\pm$ 0.5)	833.3 $\pm$ 94.3	(70.4 $\pm$ 5.2)	400.0 $\pm$ 141.4	(33.3 $\pm$ 0.0)
NUDGE	46.7 $\pm$ 18.9	(4.1 $\pm$ 0.6)	113.3 $\pm$ 81.3	(7.3 $\pm$ 3.7)	1522.2 $\pm$ 540.5	(110.0 $\pm$ 14.1)	<b>1966.7 <math>\pm</math> 237.3</b>	(101.7 $\pm$ 2.4)
BlendRL	2138.9 $\pm$ 1335.8	(15.8 $\pm$ 7.0)	<b>24755.6 <math>\pm</math> 24036.4</b>	(55.6 $\pm$ 35.1)	1955.6 $\pm$ 724.9	(101.7 $\pm$ 2.4)	1944.4 $\pm$ 245.5	(118.3 $\pm$ 22.5)
SCoBots	1055.3 $\pm$ 272.6	(- $\pm$ -)	0.0 $\pm$ 0.0	(- $\pm$ -)	2776.6 $\pm$ 1322.4	(- $\pm$ -)	0.0 $\pm$ 0.0	(- $\pm$ -)

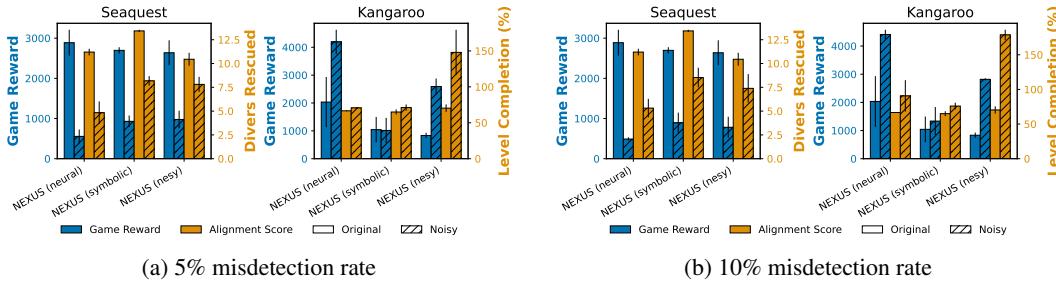
1145  
1146  
1147  
1148  
1149  
1150  
1151  
1152  
1153



(a) 5% misdetection rate (b) 10% misdetection rate

Figure 14: Misdetection and noise applied during both training and evaluation.

1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173



(a) 5% misdetection rate (b) 10% misdetection rate

Figure 15: Misdetection and noise applied only during evaluation.

1183  
1184  
1185  
1186  
1187

1188 **G LLM INTERACTION**  
11891190 We outline the procedure for leveraging an LLM to generate task-relevant skills, associated reward  
1191 functions, and a fixed meta-policy rule set. The LLM is conditioned on the game’s original manual  
1192 and structured type information describing the object-centric observations available to the agent.1193 Initially, the LLM is queried for a set of skills and corresponding reward functions. These outputs  
1194 can be manually refined before querying the LLM for a meta-policy function that selects which skill  
1195 to execute. Prior prompts and responses are retained to maintain conversational context, consistent  
1196 with standard chat behavior. [We provide the entire prompt for the game Kangaroo in Figure 16. The  
1197 prompt for Seaquest was generated equivalently.](#)  
11981199 Unedited responses from GPT-4o (via chatgpt.com on 21-07-2025) for the games Kangaroo and  
1200 Seaquest are included in Figure 17–Figure 20.1201 For Crafter, we used the LLM to generate the list of important tasks, as well as the symbolic meta-  
1202 policy function. Prompt and answers are available in Figure 21 and Figure 22. Rewards for the  
1203 specific skills were crafted manually.1204 Final implementations were modified to align with our framework constraints (e.g., JAX compatibility,  
1205 indexing conventions). [For full details, please refer to the code repository<sup>3</sup>.](#)  
1206  
1207  
1208  
1209  
1210  
1211  
1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1220  
1221  
1222  
1223  
1224  
1225  
1226  
1227  
1228  
1229  
1230  
1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238  
1239  
1240  
1241

---

 <sup>3</sup>[https://anonymous.4open.science/r/symbolic\\_options-302C/](https://anonymous.4open.science/r/symbolic_options-302C/)

1242  
1243  
1244  
1245  
1246  
1247

1248 You are a RL expert and develop a hierarchical agent to play the atari game Seaquest. Below, I have provided  
1249 a detailed description of the game. Your task is to come up with a short list of essential skills (3-4) that are  
1250 needed to solve the game.

1251 Game Description:

1252  
1253  
1254  
1255  
1256  
1257  
1258

—  
1252 <Atari game manual<sup>a</sup>>  
—

1253 Think about what the main goal of this game is, then come up with a list of the required skills to solve the  
1254 game and provide it in the output. Finally, create a pseudo reward function for each of these skills that a RL  
1255 agent can use to learn the skills.

1256 The reward functions are called at each step in the environment and retrieve the current and previous  
1257 symbolic observation of the game. Here is the type definition for the observation:  
—

1259  
1260  
1261  
1262  
1263  
1264  
1265  
1266  
1267  
1268  
1269

```
class KangarooObservation(NamedTuple):
    player_x: chex.Array
    player_y: chex.Array
    player_o: chex.Array
    platform_positions : chex.Array
    ladder_positions : chex.Array
    fruit_positions : chex.Array
    bell_position : chex.Array
    child_position : chex.Array
    falling_coco_position : chex.Array
    monkey_positions: chex.Array
    coco_positions : chex.Array
```

1270  
1271  
1272  
1273  
1274  
1275  
1276  
1277

—  
1271 Each skill and it's reward function should be as simple as possible that achieves the wanted behavior.  
1272 Provide: List of skills and description, pseudo-code reward functions for each skill.

—  
1273 <sup>a</sup>copied entirely from [https://www.atariage.com/manual\\_html\\_page.php?SoftwareLabelID=923](https://www.atariage.com/manual_html_page.php?SoftwareLabelID=923)

1278

You are now creating a function that decides at each time step in the game, which skill to choose. The skills  
1279 are<sup>a</sup>:

1280  
1281  
1282  
1283  
1284

- Obstacle Avoidance
- Vertical Navigation
- Fruit Collection
- Goal Reaching (Child Rescue)

1285 The function again has access to the current and previous observation.

1286 Think about at what point in the game you want to use which skill, then provide the function.

1287  
1288

—  
1287 <sup>a</sup>copied from answer to previous prompt, see Figure 18.

1289  
1290  
1291  
1292  
1293  
1294  
1295

Figure 16: Subsequent prompts (chat-function) to generate a skill-list and corresponding reward  
functions (top) and meta-policy function (bottom) for the game Kangaroo.

1296  
1297  
1298  
1299  
1300  
1301

### 1302 Essential Skills for Solving *Kangaroo* (Atari 2600)

1303 A hierarchical RL agent for *Kangaroo* can be built upon the following four primitive skills:

#### 1304 **1. Obstacle Avoidance**

1305 *Description:* Avoid contact with apples or monkeys via dodging or punching.

1306 *Reward Function:*

```
1308 def reward_obstacle_avoidance(obs, prev_obs):
1309     collided = any(np.array_equal(obs.player_x, x) and
1310                   np.array_equal(obs.player_y, y)
1311                   for (x, y) in obs.apple_positions + obs.monkey_positions)
1312     return -1.0 if collided else 0.0
```

#### 1313 **2. Vertical Navigation**

1314 *Description:* Climb ladders to ascend levels and progress.

1315 *Reward Function:*

```
1317 def reward_vertical_navigation(obs, prev_obs):
1318     dy = prev_obs.player_y - obs.player_y
1319     return 0.1 if dy > 0 else 0.0
```

#### 1321 **3. Fruit Collection**

1322 *Description:* Collect fruits and optionally ring a bell to regenerate them.

1323 *Reward Function:*

```
1324 def reward_fruit_collection(obs, prev_obs):
1325     reward = 0.0
1326     if len(obs.fruit_positions) < len(prev_obs.fruit_positions):
1327         reward += 0.5
1328     if np.array_equal(obs.player_x, obs.bell_position[0]) and \
1329         np.array_equal(obs.player_y, obs.bell_position[1]):
1330         reward += 0.2
1331     return reward
```

#### 1333 **4. Goal Reaching (Child Rescue)**

1334 *Description:* Reach the child's location to complete the level.

1335 *Reward Function:*

```
1336 def reward_reach_child(obs, prev_obs):
1337     return 1.0 if np.array_equal(obs.player_x, obs.child_position[0]) and \
1338                 np.array_equal(obs.player_y, obs.child_position[1]) else 0.0
```

1340 These modular rewards enable a high-level controller to invoke appropriate behaviors depending on context.

1341  
1342  
1343  
1344  
1345  
1346  
1347  
1348  
1349

Figure 17: Generated reward functions for Kangaroo by GPT4-O

1350  
1351  
1352  
1353  
1354  
1355  
1356  
1357

### Kangaroo Skill Selection Function

1358  
1359  
1360  
1361  
1362  
1363  
1364  
1365  
1366  
1367

To select appropriate skills, the agent follows a fixed priority:

1. **Obstacle Avoidance:** If any apple or monkey is in close proximity.
2. **Goal Reaching:** If the player is near the child on the top platform.
3. **Fruit Collection:** If fruit is visible and reachable.
4. **Vertical Navigation:** Default behavior to explore or ascend.

### Skill Selector Function:

```
1368 def select_skill (obs, prev_obs):
1369     player_pos = np.array([obs.player_x, obs.player_y])
1370
1371     def in_danger():
1372         danger_entities = obs.apple_positions + obs.monkey_positions
1373         return any(np.linalg.norm(player_pos - np.array([x, y])) <= 1.0
1374             for (x, y) in danger_entities)
1375
1376     def near_child():
1377         child_pos = np.array([obs.child_position[0], obs.child_position[1]])
1378         return np.linalg.norm(player_pos - child_pos) <= 1.5
1379
1380     def fruit_visible():
1381         return any(abs(y - obs.player_y) <= 1 and
1382                 abs(x - obs.player_x) <= 3
1383                 for (x, y) in obs.fruit_positions)
1384
1385     if in_danger():
1386         return obstacle_avoidance
1387     elif near_child():
1388         return goal_reaching
1389     elif fruit_visible():
1390         return fruit_collection
1391     else:
1392         return vertical_navigation
```

This rule-based selection policy balances safety, progression, and reward acquisition, enabling effective hierarchical control.

1392  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403

Figure 18: Generated fixed meta-policy function for Kangaroo by GPT4-O

1404  
1405Essential Skills for Solving *Seaquest* (Atari 2600)1406  
1407

To solve *Seaquest* effectively, a hierarchical RL agent must master the following four core skills:

1408

**1. Avoid Obstacles (Survival Navigation)**

1409

*Description:* Avoid collisions with sharks, enemy submarines, missiles, and the patrol sub, both underwater and when surfacing.

1410

*Reward Function (pseudo-code):*

1411

```
1413 def avoid_obstacles_reward (prev_obs, curr_obs):
1414     if player_collided (curr_obs):
1415         return -1.0
1416     return 0.0
1417
1418 def player_collided (obs):
1419     return check_collision_with_sharks (obs) or \
1420         check_collision_with_subs (obs) or \
1421         check_collision_with_missiles (obs) or \
1422         check_collision_with_surface_sub (obs)
```

1412

**2. Collect Divers**

1413

*Description:* Navigate to and collect divers to accumulate score and progress.

1414

*Reward Function (pseudo-code):*

1415

```
1416 def collect_divers_reward (prev_obs, curr_obs):
1417     prev_count = count_nearby_divers (prev_obs)
1418     curr_count = count_nearby_divers (curr_obs)
1419     return float (prev_count > curr_count) * 1.0
1420
1421 def count_nearby_divers (obs):
1422     return sum (is_near (obs.player_x, obs.player_y, dx, dy)
1423                 for dx, dy in obs.diver_positions )
```

1421

1422

1423

1424

1425

**3. Surface Strategically**

1426

*Description:* Surface safely when carrying six divers or when oxygen is low, to drop off divers and replenish oxygen.

1427

*Reward Function (pseudo-code):*

1428

```
1429 def surface_reward (prev_obs, curr_obs):
1430     surfaced_now = prev_obs.player_y > 0 and curr_obs.player_y == 0
1431     has_6_divers = num_divers_collected (prev_obs) == 6
1432     low_oxygen = prev_obs.oxygen < OXYGEN_THRESHOLD
1433     if surfaced_now and (has_6_divers or low_oxygen):
1434         return 2.0 if has_6_divers else 0.5
1435     return 0.0
1436
1437 def num_divers_collected (obs):
1438     return 6 - sum (is_valid (dx, dy) for dx, dy in obs.diver_positions )
```

1436

1437

1438

**4. Eliminate Threats (Combat)**

1439

*Description:* Fire missiles to destroy sharks and enemy submarines, clearing paths and increasing score.

1440

*Reward Function (pseudo-code):*

1441

```
1442 def eliminate_threats_reward (prev_obs, curr_obs):
1443     num_enemies_prev = count_enemies (prev_obs)
1444     num_enemies_curr = count_enemies (curr_obs)
1445     return float (num_enemies_curr < num_enemies_prev) * 0.5
1446
1447 def count_enemies (obs):
1448     return sum (is_valid (dx, dy) for dx, dy in obs.shark_positions ) + \
1449                 sum (is_valid (dx, dy) for dx, dy in obs.sub_positions )
```

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458

1459

## 1460 Seaquest Skill Selection Function

1461

1462 Here is a skill selection function that chooses the appropriate control skill at each time step  
1463 based on game context.

1464

## Skill Selection Function (pseudo-code):

1465

1466

1467

1468

1469

1470

1471

1472

1473

1474

1475

1476

1477

1478

1479

1480

1481

1482

1483

1484

1485

1486

1487

1488

1489

1490

1491

1492

1493

1494

1495

1496

1497

1498

1499

```

def select_skill (prev_obs, curr_obs):
    # Skill names
    AVOID = "obstacle_avoidance"
    COLLECT = "collect_divers"
    SURFACE = "surface_strategically"
    COMBAT = "eliminate_threats"

    # === Utility Functions ===
    def is_threat_near (obs: SeaquestObservation) -> bool:
        return check_collision_with_sharks (obs) or \
               check_collision_with_subs (obs) or \
               check_collision_with_missiles (obs) or \
               check_collision_with_surface_sub (obs)

    def num_divers_collected (obs: SeaquestObservation) -> int:
        return 6 - sum(is_valid (x, y) for x, y in obs.diver_positions)

    def low_oxygen(obs: SeaquestObservation) -> bool:
        return obs.oxygen < 20 # adjustable threshold

    def at_surface (obs: SeaquestObservation) -> bool:
        return obs.player_y == 0

    # === Skill Selection Logic ===
    if is_threat_near (curr_obs):
        return AVOID

    if (num_divers_collected (curr_obs) == 6 or low_oxygen(curr_obs)) and \
        not at_surface (curr_obs):
        return SURFACE

    if has_nearby_enemies(curr_obs): # e.g. visible on screen in front
        return COMBAT

    return COLLECT

```

## Explanation of Logic:

1500

1501

1502

1503

1504

1505

1506

1507

1508

1509

1510

1511

Figure 20: Generated fixed meta-policy for Seaquest by GPT4-O

1512  
 1513 You are a reinforcement learning specialist. The following is the basic description of the game Crafter, which  
 1514 you are trying to solve.  
 1515  
 1516 Crafter is a game about exploring dungeons, mining, crafting and fighting enemies. The player can move in  
 1517 the four cardinal directions using WASD and can interact using SPACE. Interacting can cause the player to  
 1518 attempt to mine (a block), attack (a creature), drink (water or from a fountain), or eat (fruit).  
 1519 The player has the 'intrinsic' health, hunger, thirst and energy. Hunger, thirst and energy will naturally  
 1520 decrease and must be replenished by eating, drinking and sleeping respectively. Health will recover when  
 1521 hunger, thirst and energy are non-zero and will decrease if any of these are 0. If the player's health falls  
 1522 beneath 0 they will die and the game will restart.  
 1523  
 1524 First, come up with 5 basic skills that can be used to solve the game. Only provide necessary skills without  
 1525 duplications. After that, please come up with a simple function (pseudo code) that gets the game-state as  
 1526 input and returns an integer deciding on the active skill to select in the current state. The skills correspond to  
 1527 the ones that you chose before. Try to solve the game as far as possible.  
 1528 The game-state looks like this:  
 1529  
 1530 

```
@struct, dataclass
1531 class Inventory:
1532     wood: int = 0
1533     stone: int = 0
1534     coal: int = 0
1535     iron: int = 0
1536     diamond: int = 0
1537     sapling: int = 0
1538     wood_pickaxe: int = 0
1539     stone_pickaxe: int = 0
1540     iron_pickaxe: int = 0
1541     wood_sword: int = 0
1542     stone_sword: int = 0
1543     iron_sword: int = 0
```

 1544  
 1545 

```
@struct, dataclass
1546 class Mobs:
1547     position: jnp.ndarray
1548     health: int
1549     mask: bool
1550     attack_cooldown: int
```

 1551  
 1552 

```
@struct, dataclass
1553 class EnvState:
1554     map: jnp.ndarray
1555     mob_map: jnp.ndarray
1556     player_position: jnp.ndarray
1557     player_direction: int
1558     # Intrinsic
1559     player_health: int
1560     player_food: int
1561     player_drink: int
1562     player_energy: int
1563     is_sleeping: bool
1564     # Second order intrinsic
1565     player_recover: float
1566     player_hunger: float
1567     player_thirst: float
1568     player_fatigue: float
1569     inventory: Inventory
1570     zombies: Mobs
1571     cows: Mobs
1572     skeletons: Mobs
1573     arrows: Mobs
1574     arrow_directions: jnp.ndarray
1575     growing_plants_positions: jnp.ndarray
1576     growing_plants_age: jnp.ndarray
1577     growing_plants_mask: jnp.ndarray
```

1566  
1567  
1568  
1569  
1570  
1571

### Essential Skills for Solving *Crafter*

1572  
1573  
1574

**Basic Skills.** We identify the following five skills as sufficient building blocks for solving *Crafter*:

1575  
1576  
1577  
1578  
1579  
1580  
1581  
1582  
1583  
1584  
1585

1. **Gather Resources:** Mine wood, stone, coal, iron, or other materials when available.
2. **Craft Tools/Weapons:** Use gathered resources to craft pickaxes and swords for efficiency and combat readiness.
3. **Eat/Drink/Sleep:** Replenish food, water, or energy to maintain survival.
4. **Fight/Defend:** Attack hostile mobs when they are nearby and health is sufficient.
5. **Explore/Move:** Traverse unexplored areas to discover resources, water, and safe locations.

1586  
1587  
1588  
1589  
1590  
1591  
1592  
1593  
1594  
1595  
1596  
1597  
1598  
1599  
1600  
1601  
1602  
1603  
1604  
1605  
1606  
1607  
1608  
1609  
1610  
1611  
1612

**Policy Function.** The following pseudo-code outlines a rule-based strategy for selecting the active skill based on the current game state.

```
def select_skill(state: EnvState) -> int:
    """
    Returns the skill index:
    0 = Gather Resources
    1 = Craft Tools/Weapons
    2 = Eat/Drink/Sleep
    3 = Fight/Defend
    4 = Explore/Move
    """

    # --- Survival priority ---
    if state.player_food <= 2 or state.player_drink <= 2 or state.
        player_energy <= 2:
        return 2 # Eat/Drink/Sleep

    # --- Combat ---
    if mob_nearby(state) and state.player_health > 3:
        return 3 # Fight/Defend

    # --- Resource management ---
    if resources_nearby(state) and inventory_needs_tools(state.
        inventory):
        return 0 # Gather Resources

    if can_craft_better_tools(state.inventory):
        return 1 # Craft Tools/Weapons

    # --- Default action ---
    return 4
```

1613  
1614  
1615

Figure 22: Generated list of essential skills and a rule-based policy for solving *Crafter* by GPT4-O. The policy prioritizes survival, then combat, followed by resource gathering and crafting, and defaults to exploration otherwise.

1616  
1617  
1618  
1619