
Sliding Puzzles Gym: A Scalable Benchmark for State Representation in Visual Reinforcement Learning

Bryan L. M. de Oliveira^{1,2} Luana G. B. Martins¹ Bruno Brandão^{1,2} Murilo L. da Luz^{1,2}
Telma W. de L. Soares^{1,2} Luckeciano C. Melo^{1,3}

Abstract

Effective visual representation learning is crucial for reinforcement learning (RL) agents to extract task-relevant information from raw sensory inputs and generalize across diverse environments. However, existing RL benchmarks lack the ability to systematically evaluate representation learning capabilities in isolation from other learning challenges. To address this gap, we introduce the Sliding Puzzles Gym (SPGym), a novel benchmark that transforms the classic 8-tile puzzle into a visual RL task with images drawn from arbitrarily large datasets. SPGym’s key innovation lies in its ability to precisely control representation learning complexity through adjustable grid sizes and image pools, while maintaining fixed environment dynamics, observation, and action spaces. This design enables researchers to isolate and scale the visual representation challenge independently of other learning components. Through extensive experiments with model-free and model-based RL algorithms, we uncover fundamental limitations in current methods’ ability to handle visual diversity. As we increase the pool of possible images, all algorithms exhibit in- and out-of-distribution performance degradation, with sophisticated representation learning techniques often underperforming simpler approaches like data augmentation. These findings highlight critical gaps in visual representation learning for RL and establish SPGym as a valuable tool for driving progress in robust, generalizable decision-making systems.

¹Advanced Knowledge Center for Immersive Technologies – AKCIT, Brazil ²Institute of Informatics, Federal University of Goiás, Goiânia, Brazil ³OATML, University of Oxford. Correspondence to: Bryan L. M. de Oliveira <bryanlincoln@discente.ufg.br>.

1. Introduction

Learning meaningful representations from raw sensory inputs, such as visual data, is fundamental to reinforcement learning (RL) agents’ ability to generalize across different tasks in complex, open-world environments (Bengio et al., 2013; Lesort et al., 2018). In visual RL, agents must process high-dimensional pixel data, extract useful features, and utilize these features for decision-making (Mnih et al., 2015; Yarats et al., 2021a). This becomes especially crucial as real-world applications demand adaptability to unstructured and diverse observations. However, measuring an agent’s representation learning capabilities independently of other learning tasks, such as policy optimization or dynamics modeling, remains a key challenge in RL benchmarks.

While traditional RL benchmarks such as Atari (Bellemare et al., 2013) and DeepMind Control Suite (Tassa et al., 2018) provide valuable testing grounds for overall agent performance, they inherently intertwine representation learning with policy optimization and environment dynamics. More recent specialized benchmarks for visual learning, though promising, fall short in their ability to precisely modulate visual complexity. ProcGen (Cobbe et al., 2020) concurrently alters both visual and task difficulty, obscuring the specific impact of representation learning, while the Distracting Control Suite (Stone et al., 2021) incorporates visual variations that are orthogonal to the core task objectives. Consequently, existing benchmarks lack the necessary precision to systematically assess an agent’s capacity to acquire task-relevant visual representations in isolation.

To address this critical gap in visual RL evaluation, we introduce the *Sliding Puzzles Gym (SPGym)*¹, an open-source benchmark specifically designed to assess how agents handle increasingly diverse visual observations. As depicted in Figure 1, SPGym transforms the classic 8-tile puzzle into a visual RL task through three innovative design principles: (1) maintaining consistent environment dynamics regardless of difficulty level, ensuring that the underlying task remains fixed; (2) enabling precise control over visual complexity

¹Available at <https://github.com/bryanoliveira/sliding-puzzles-gym>.

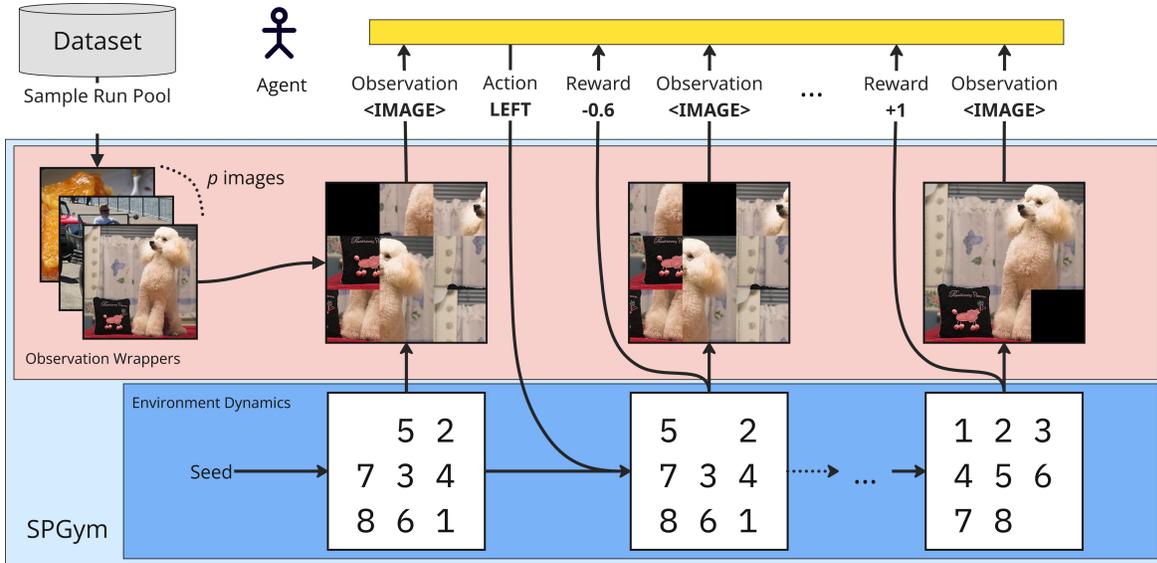


Figure 1. **Overview of SPGym.** The framework extends the 8-tile puzzle by replacing numbered tiles with image patches. At each training run, SPGym samples a pool of images and, at each episode, it randomly selects one of those images to form the observations. While we scale visual diversity by adjusting the pool size, the task and environment dynamics remain fixed.

through adjustable grid dimensions and image pool sizes; and (3) establishing an unambiguous success metric based on puzzle completion. This carefully engineered framework provides researchers with a controlled environment to systematically evaluate how agents’ performance degrades with increasing visual diversity, offering insights into the fundamental limitations of their representation learning capabilities while keeping the core task unchanged.

Our experiments with SPGym reveal that it effectively distinguishes agents by their representation learning capabilities, showing a strong correlation with task performance. While pretrained encoders and data augmentation provide clear benefits, many advanced methods surprisingly underperform their baselines, suggesting their assumptions don’t transfer to SPGym’s unique blend of visual complexity and structured dynamics. More concerningly, our tiered generalization analysis exposes fundamental limitations: agents achieve perfect performance on training images but fail to transfer to unseen ones, even with larger training pools. Performance often degrades as pool size increases, highlighting a key trade-off: greater visual diversity increases learning difficulty without improving generalization, suggesting current methods struggle to learn truly generalizable representations. We further demonstrate SPGym’s extensibility through experiments with procedurally generated images and larger puzzles, maintaining fixed observation/action spaces while expanding state space and visual diversity. These findings expose critical gaps in current visual RL methods and establish SPGym as a valuable tool for advancing robust, generalizable decision-making systems.

Contributions. We make three key contributions to visual RL research: (1) SPGym, a novel benchmark that systematically evaluates representation learning by scaling visual complexity while keeping environment dynamics constant; (2) an extensive empirical analysis of state-of-the-art methods, uncovering critical limitations in their ability to process diverse visual inputs; and (3) fundamental insights into the challenges of scaling visual RL, offering directions for advancing representation learning in decision-making systems.

2. Related Work

Traditional RL benchmarks. Foundational visual RL benchmarks like the Atari Learning Environment (Belle-mare et al., 2013), DeepMind Control Suite (Tassa et al., 2018), DeepMind Lab (Beattie et al., 2016), and CARLA (Dosovitskiy et al., 2017) pioneered pixel-based agent evaluation. These environments drove advances in representation learning techniques including pretraining (Higgins et al., 2017; Stooke et al., 2021; Schwarzer et al., 2021), contrastive learning (Laskin et al., 2020b), self-supervised prediction (Schwarzer et al., 2020), and world models (Hafner et al., 2025). However, by assessing overall agent performance across representation learning, policy optimization, and dynamics modeling, these benchmarks obscure the specific impact of representation learning improvements.

SPGym addresses this limitation by precisely controlling visual complexity while keeping environment dynamics fixed, enabling focused evaluation of representation learning.

Specialized benchmarks for visual RL. Recent benchmarks have sought to better assess visual generalization. ProcGen (Cobbe et al., 2020) employs procedurally generated levels, but its difficulty scaling simultaneously affects both visual and task complexity. The Distracting Control Suite (Stone et al., 2021) evaluates robustness to visual variations, but these serve as task-irrelevant distractions that agents can safely ignore while solving core tasks. COOM (Tomilin et al., 2023) provides a continual learning benchmark for embodied pixel-based RL in 3D environments, focusing on catastrophic forgetting and knowledge transfer across tasks rather than isolating representation learning in a single task. These approaches fail to create representation learning challenges essential for task success.

SPGym overcomes these limitations by making visual understanding fundamental to puzzle solving while maintaining fixed task complexity. Through controlled scaling of visual diversity in puzzle tiles, SPGym creates representation learning challenges tied to task completion, enabling precise evaluation of agents’ visual learning capabilities.

Puzzle-based benchmarks. Estermann et al. (2024) demonstrated puzzle-based environments’ value for evaluating neural algorithmic reasoning, focusing on discrete states. Though they tested pixel observations with sparse rewards, agents struggled with basic visual inputs. This work revealed puzzles’ potential and challenges for visual learning, showing the need for systematic representation learning evaluation in these controlled settings.

SPGym advances this research direction by incorporating rich visual observations while maintaining the controlled nature of puzzle environments, facilitating systematic assessment of representation learning in the visual domain.

Methods for solving the sliding tile puzzle. Classical approaches to sliding puzzles employ heuristic search algorithms like A* and IDA* with domain-specific heuristics such as Manhattan distance (Korf, 1985; Burns et al., 2012; Lee & See, 2022). While these methods guarantee optimal solutions, they require direct access to internal puzzle states and are computationally demanding. Deep RL provides a scalable alternative that learns effective strategies without hand-crafted heuristics (Agostinelli et al., 2019; Moon & Cho, 2024; Estermann et al., 2024). However, prior work has predominantly focused on discrete state representations rather than learning from visual observations (Agostinelli et al., 2019; Moon & Cho, 2024).

Building on this foundation, our work evaluates standard RL algorithms on SPGym, where agents must learn exclusively from pixel observations without access to internal states, offering insights into their out-of-the-box performance in this challenging visual learning setting.

3. The Sliding Puzzles Gym

SPGym extends the classic sliding puzzle, a game where players rearrange shuffled tiles on a grid by sliding a tile into an adjacent empty space, with the goal of restoring an ordered configuration (Figure 1). Our framework generalizes this by supporting configurable $H \times W$ grid dimensions (from 2×2 upwards) and diverse observation modalities. In our experiments, we primarily use a 3×3 grid where tiles are image patches, and the agent receives the composite image of the grid as input. SPGym adheres to the Gym (Brockman, 2016) interface, promoting modularity between environment and agent.

Formalization. We formulate SPGym as a partially observable Markov decision process (POMDP) defined by the tuple $(\mathcal{S}, \mathcal{X}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \mathcal{S}_0)$, where \mathcal{S} is the finite state space representing the internal configuration of puzzle tiles, \mathcal{X} is the observation space of image-based tiles, \mathcal{A} is the action space, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ defines the deterministic transition dynamics, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, and \mathcal{S}_0 is the initial state distribution.

State space and observations. Crucially, agents do not have direct access to states $s \in \mathcal{S}$ and must learn policies $\pi : \mathcal{X} \rightarrow \mathcal{A}$ based solely on observations $x \in \mathcal{X}$. For visual observations, each training run begins by sampling p images from a predefined dataset to form an image pool \mathcal{I} . At the start of each episode, we select a random image $i \in \mathcal{I}$ and split it into $H \times W$ distinct, spatially indexed patches, and overlay it onto the puzzle. The agent’s task is therefore to reconstruct the shuffled image, testing its ability to form compositional representations from pixels. This formulation provides two independent mechanisms for controlling complexity: (1) varying p adjusts the diversity of the observation space \mathcal{X} by changing the pool of available images, and (2) modifying H and W alters the state space complexity by changing the grid dimensions and number of puzzle pieces. Both mechanisms operate while keeping the underlying dynamics \mathcal{P} , the action space \mathcal{A} , and the reward function \mathcal{R} fixed. SPGym is dataset-agnostic, supporting any image dataset, including procedurally generated ones.

Action space and dynamics. The action space \mathcal{A} contains four possible actions: *UP*, *DOWN*, *LEFT*, or *RIGHT*, which move the corresponding tile to the empty space in a discrete manner. Once the agent selects a tile to move, the dynamics \mathcal{P} define the next puzzle state in a predictable and deterministic way, as illustrated in Figure 1. The objective is to rearrange shuffled image tiles into their correct positions, forming a complete, intelligible image.

Reward function. The reward function \mathcal{R} provides dense feedback r to guide agent learning while accommodating

the non-monotonic nature of optimal puzzle-solving paths. Following established approaches in the sliding puzzle literature (Korf, 1985; Burns et al., 2012; Lee & See, 2022; Moon & Cho, 2024), we base our reward on the Manhattan distance between each tile’s current position and its target position, normalized across all tiles. This distance metric is widely adopted in computational solutions to sliding puzzles due to its efficiency and effectiveness in capturing progress toward the goal state. Specifically, at each time step, the reward is computed as follows:

$$r_t = \begin{cases} -D, & \text{if action is valid} \\ -1, & \text{if action is invalid} \\ +1, & \text{if puzzle is solved} \end{cases}, \text{ with} \quad (1)$$

$$D = \frac{\sum_{i=1}^H \sum_{j=1}^W |u_{i,j} - u_{i,j}^*| + |v_{i,j} - v_{i,j}^*|}{\sum_{i=1}^H \sum_{j=1}^W \max(i, H - i) + \max(j, W - j)}. \quad (2)$$

Here, $(u_{i,j}, v_{i,j})$ represents the current position of the tile at index (i, j) , $(u_{i,j}^*, v_{i,j}^*)$ is its target position. D is the normalized Manhattan distance between current and target positions. Invalid actions, such as attempting to move a tile outside the grid boundaries, result in a penalty of -1 and don’t alter the puzzle state. For example, in Figure 1, the *DOWN* and *RIGHT* actions would be invalid for the first state. Successfully solving the puzzle rewards the agent with $+1$ and terminates the episode. This formulation provides a well-behaved learning signal between $[-1, +1]$ and encourages solving the puzzle in minimal steps.

Initial state distribution. SPGym’s initial state distribution \mathcal{S}_0 can be defined through two methods. The primary method generates a uniformly random $H \times W$ array where all tiles (including the blank tile) are placed randomly, and ensures solvability by adjusting the puzzle’s parity through swapping the first two tiles when needed (Johnson & Story, 1879). The second method, designed to support curriculum learning, starts from a solved state and applies a sequence of random valid moves to reach the initial configuration. Although this approach guarantees solvability due to the reversibility of moves, it becomes computationally expensive for larger grid sizes due to the sequential nature of move application. Given these efficiency considerations, we exclusively employ the first method in our experiments.

Diversity scalability and extensibility. While our main experiments focus on 3×3 grids with image observations, SPGym supports variable grid sizes from 2×2 to larger $H \times W$ configurations. We demonstrate this with experiments on 4×4 grids (Table 3), which significantly increase sample complexity compared to 3×3 grids. We also use one-hot observation spaces, illustrated in Figure 2, as a

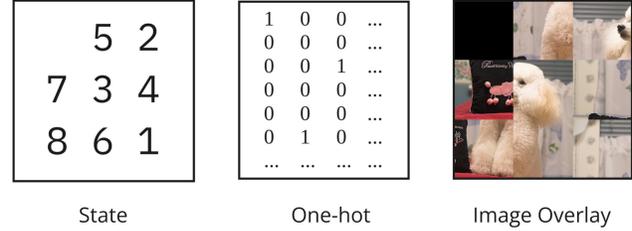


Figure 2. **Different observation modalities in SPGym.** Each modality presents a unique challenge for representation learning. The three presented observations represent the same puzzle state. We focus our experiments on image-based observations.

proxy for ground-truth puzzle states to benchmark the visual representation challenge (results in Appendix C.1).

The primary scaling mechanism in SPGym involves increasing visual diversity by expanding the image pool while keeping the grid size fixed (Figure 3). As shown in Table 1, larger image pools consistently increase sample complexity. Since core task elements like state and action spaces, transition dynamics, and reward function remain unchanged, this increased difficulty primarily reflects the challenge of learning robust visual representations. The quality of learned representations is directly linked to task performance, as demonstrated by linear probe analyses (Appendix C.1.2). Probe accuracy decreases with larger pool sizes (Table 10), exhibiting strong correlation with sample efficiency. Comparisons with one-hot encoded baselines (Appendix C.1.1), where the representation learning burden is minimal, further highlight this focus. While the one-hot setting offers fixed difficulty, the systematic increase in sample complexity for image-based agents as visual diversity scales provides a stress test not possible in the one-hot setting.

The secondary scaling mechanism adjusts grid dimensions. Increasing from 3×3 to 4×4 dramatically increases the number of steps to solve the task due to expanded state space complexity. While 3×3 puzzles have approximately 1.81×10^5 possible states, 4×4 puzzles present a dramatically larger challenge with approximately 10^{13} possible states. This increase in puzzle size not only expands the state space exponentially but also makes the representation learning problem harder, as the same image must now be divided into more patches that need to be correctly positioned. Importantly, larger grids don’t affect the observation space \mathcal{X} (image sizes remain fixed) or action space \mathcal{A} (still four directional moves), and both puzzle sizes maintain identical observation and action spaces, isolating the impact of state space complexity. While this mechanism evaluates exploration and policy-learning capabilities, our experiments demonstrate that 3×3 grids provide sufficient discriminative power for visual representation learning assessment.

By maintaining consistent environment dynamics across

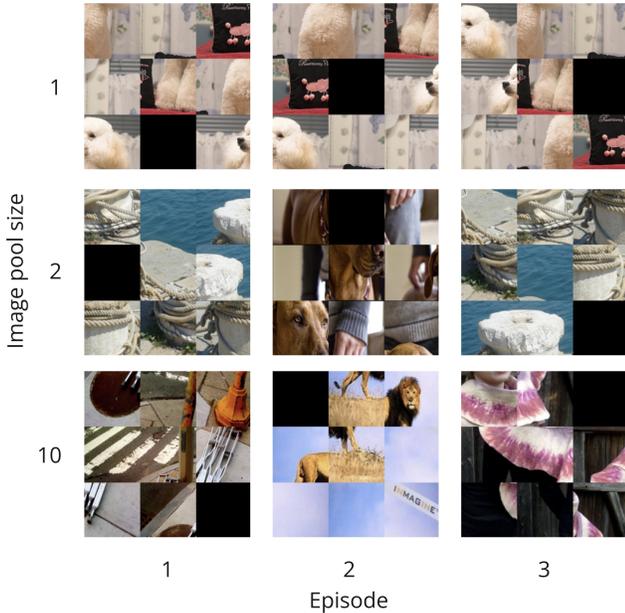


Figure 3. **The visual diversity scalability of SPGym.** Each row presents the first observation of 3 different episodes of the same training run. The visual diversity scales with the size of the image pool. Crucially, we keep the grid size fixed, ensuring that the increased difficulty comes solely from the agent’s need to handle a larger variety of visual observations.

these scaling mechanisms, SPGym creates a controlled experimental setting where performance variations can be more directly attributed to an agent’s representation learning and generalization capabilities.

4. Methods

In this section, we detail the evaluation protocol for various RL agents within the SPGym framework, focusing on their ability to manage increasing visual diversity.

Experimental setup. We sample images from ImageNet-1k’s validation split (Russakovsky et al., 2015) to construct visual observations, resizing each image to 84×84 pixels and normalizing values to $[0, 1]$. To isolate visual diversity, we fix the puzzle size to 3×3 and vary only the image pool size p . For each training run, we randomly sample p distinct images to create a fixed image pool. At episode start, we randomly select one image from this pool to generate the puzzle observations. We cap the number of environment steps to 10M and limit episodes to 1,000 steps.

Our analysis focuses on sample efficiency, measured by the number of environment steps required to solve the puzzle (lower is better). We calculate this metric by averaging the steps needed to reach 80% success rate across all parallel environments in a run, then average this number across seeds.

We terminate training runs early when an agent maintains 100% success rate for 100 consecutive episodes, indicating task completion. This early termination serves two purposes: it enables out-of-distribution evaluation before extreme encoder overfitting occurs, and it reduces computational costs for running the comprehensive set of experiments. Each experiment is conducted 5 times with different random seeds, and we report the mean ± 1.96 standard errors (95% confidence interval).

We evaluate sample efficiency using pools of 1, 5, and 10 images across all agent variants. For standard agents, we additionally test with progressively larger pools up to 100 images or until final performance is less than 80% success rate after the full 10M training steps. This protocol allows us to analyze both the effectiveness and scalability of different algorithms and representation learning methods.

Algorithms and variants. We explore three distinct algorithmic approaches: Soft Actor-Critic (SAC) (Haarnoja et al., 2018a), Proximal Policy Optimization (PPO) (Schulman et al., 2017), and DreamerV3 (Hafner et al., 2025), each representing different strategies for learning from visual observations. For each algorithm, we evaluate several representation learning variants to assess their effectiveness in handling visual diversity.

For SAC, we begin with the standard implementation for discrete action spaces proposed by Christodoulou (2019). We then examine several representation learning variants: RAD (Laskin et al., 2020a), which employs data augmentation; CURL (Laskin et al., 2020b), which uses contrastive learning; SPR (Schwarzer et al., 2020), which incorporates self-supervised prediction; DBC (Zhang et al., 2021), which focuses on state metric learning; SAC-AE and SAC-VAE (Yarats et al., 2021b), which utilize reconstruction-based learning; and Simple Baseline (SB) (Tomar et al., 2023), which implements a simplified approach with reward and transition prediction.

For PPO, we evaluate three encoder configurations: the standard version with random initialization, a variant pretrained on the same image distribution (in-distribution, ID), and a variant pretrained on a different image distribution (out-of-distribution, OOD). These pretrained variants respectively provide an upper bound on expected pretraining performance and help establish the potential benefits of generally pretrained encoders, though we acknowledge this may make direct sample efficiency comparisons with other methods less fair.

For DreamerV3, we compare the standard version against a variant without decoder gradients to evaluate the impact of the reconstruction objective on performance. Detailed descriptions for all agents and their learning objectives are provided in Appendix B.

Hyperparameters. RAD, CURL, and SPR require data augmentation. We apply these augmentations to observations after sampling them from the replay buffer. For each algorithm, we conducted individual augmentation searches with the objective of maximizing sample efficiency (detailed in Appendix A.4.2). These experiments consistently converged to a simple two-step pipeline, which we use throughout all evaluations: first converting to grayscale with 20% probability, then randomly shuffling the color channels.

We evaluate out-of-the-box performance of existing approaches in SPGym by adopting neural architectures and hyperparameters from established visual discrete control implementations. Our base architecture uses three-layer CNN encoders with mirrored deconvolutional decoders, while actor, critic, and auxiliary components employ multi-layer perceptrons (MLPs).

For SAC-based agents, we follow Yarats et al. (2021b) and Tomar et al. (2023) by blocking actor gradients through the encoder while allowing critic and auxiliary gradients to prevent representation collapse. PPO and DreamerV3 maintain their original gradient flow patterns. While PPO and DreamerV3 worked robustly with default configurations, SAC-based agents required tuning of the temperature parameter α (Appendix A.4.1). We found a fixed value of 0.05 to work best across all SAC variants, as the automatic tuning from Haarnoja et al. (2018b) proved ineffective here.

Where applicable, we preserve uniform hyperparameter while drawing algorithm-specific configurations from their respective source papers. The publicly available code repository² extends both CleanRL (Huang et al., 2022) and the official DreamerV3 codebase. Additional methodological details are provided in Appendix A.

5. Results

Our analysis reveals three fundamental tensions in visual RL: between method assumptions and environment structure, sample efficiency and solution optimality, and training diversity versus generalization capability. We organize findings through seven research questions.

5.1. Can SPGym distinguish agents on representation learning capabilities?

Table 1 demonstrates SPGym’s ability to differentiate agents based on their representation learning approaches across varying levels of visual diversity. Our primary goal in this analysis is to compare the sample efficiency with which different agents learn useful visual representations. Notably, certain representation learning methods provide clear

²Accessible at <https://github.com/bryanoliveira/spgym-experiments>

Table 1. Million steps to reach 80% success rate across pool sizes. Lower is better. Best performing variant for each algorithm and pool size is highlighted in bold.

Agent	Pool 1	Pool 5	Pool 10
PPO	1.75±0.44	7.80±1.08	9.73±0.36
PPO + PT (ID)	0.95±0.21	5.55±1.22	9.17±1.10
PPO + PT (OOD)	1.34±0.42	7.03±1.07	9.70±0.41
SAC	0.33±0.07	0.91±0.12	2.03±0.38
SAC + RAD	0.24±0.03	0.42±0.06	0.82±0.18
SAC + CURL	0.46±0.10	1.56±0.31	5.24±1.92
SAC + SPR	2.09±0.81	3.68±1.68	10.00±0.00
SAC + DBC	0.99±0.25	1.12±0.22	2.13±0.41
SAC + AE	1.04±0.24	1.02±0.19	2.01±0.38
SAC + VAE	1.13±0.14	5.30±0.68	10.00±0.00
SAC + SB	0.98±0.88	2.08±0.30	10.00±0.00
DreamerV3	0.42±0.06	1.23±0.20	1.44±0.58
DreamerV3 _{w/o dec.}	1.13±0.12	1.79±0.61	2.57±0.91

benefits, though the extent can depend on the pool size.

To explore the potential benefits of leveraging pretrained encoders in this setup, we evaluated PPO with two types of pretrained encoders. It is important to acknowledge that direct comparisons of sample efficiency with encoders trained from scratch can be nuanced, as pretrained encoders have inherently been exposed to significantly more data during their pretraining phase. The in-distribution pretrained encoder (PT (ID)) for PPO, intended as a proxy to estimate an approximate upper bound on performance achievable with pretraining, significantly boosts sample efficiency across all tested pool sizes, almost matching the performance of PPO with one-hot observations (see Appendix C.1.1). The out-of-distribution pretrained encoder (PT (OOD)), designed to represent the average performance one might expect from a generally pretrained encoder, offers more modest gains which also decrease with larger pools.

For SAC, data augmentation via RAD consistently and substantially improves efficiency across all tested pool sizes, demonstrating increasing benefits as visual diversity scales. The improvements are particularly pronounced for larger pools, making SAC significantly faster at handling visual diversity. Conversely, many sophisticated auxiliary representation learning methods struggle in SPGym. For instance, CURL, SPR, and VAE variants consistently require more samples than standard SAC across all tested pool sizes, with particularly poor performance on larger pools. DBC also generally underperforms or offers marginal improvements compared to standard SAC.

DreamerV3 demonstrates particularly strong performance, consistently outperforming both PPO and SAC variants across all pool sizes. The full DreamerV3 model shows remarkably stable performance across different pool sizes, sig-

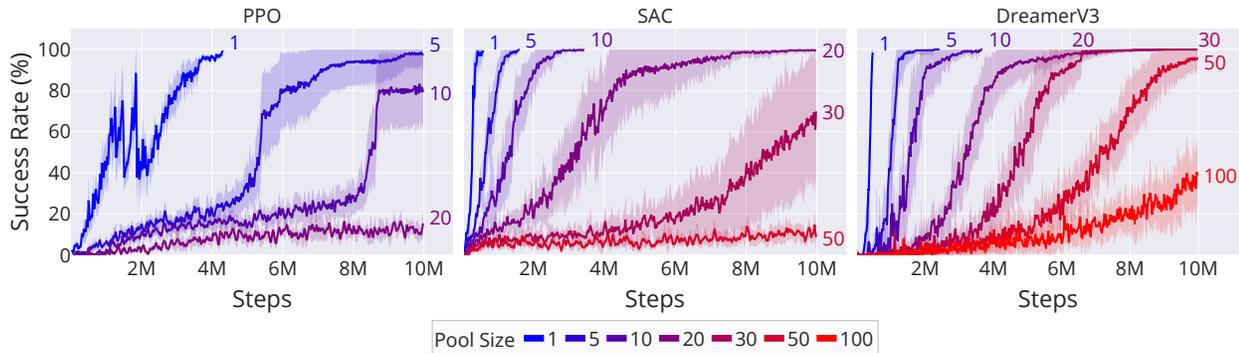


Figure 4. **Success rate as a function of environment steps.** The gradual increase in visual diversity affects the sample efficiency of standard PPO, SAC, and DreamerV3 agents at different rates. Each line represents a different pool size, from 1 to 100 images.

nificantly better than most other methods. The variant without decoder gradients shows reduced performance across all configurations, highlighting the importance of the world model’s discrete reconstruction objective. This suggests that DreamerV3’s approach of learning a predictive model of the environment provides a particularly effective foundation for handling visual diversity in SPGym.

These results highlight SPGym’s value as a diagnostic tool, effectively distinguishing which approaches can handle its distinct combination of visual diversity and structured dynamics. The benchmark exposes important limitations in methods that have proven successful in other visual RL domains, suggesting that their underlying assumptions may not align with SPGym’s unique characteristics. For detailed performance analysis and hypothesized explanations for agents behavior, see Appendix C.2.1.

5.2. How does visual diversity affect agents?

Figure 4 exposes critical limitations in current methods. While all agents eventually degrade with larger image pools, their failure modes differ. PPO degrades significantly at pool size 10 and fails at 20. SAC performs well up to pool size 20, degrades at 30, and fails at 50. DreamerV3 shows the most robust scaling, learning effectively at pool size 50 and exhibiting learning even at pool size 100.

These distinct degradation patterns suggest differences in handling increasing visual diversity, potentially due to agents memorizing features instead of learning generalizable representations, thus exhausting network capacity. PPO’s early failure indicates difficulty forming robust representations and rapid capacity exhaustion. SAC’s better performance on medium pools, possibly aided by its replay buffer and encoder training independent of actor gradients, suggests stronger initial representation learning. However, it also faces fundamental limitations, potentially as memo-

rization becomes unsustainable. DreamerV3’s world model architecture appears to foster more compressed representations, leading to more graceful degradation, but it too struggles with very large pools.

Preliminary DreamerV3 experiments with extremely large image pools (10,000-20,000 images) support this interpretation: agents failed to learn useful policies as each observation became virtually unique. This suggests that increased dataset scale alone is insufficient if the RL signal is too sparse for encoder training amidst high diversity. These methods appear to memorize rather than generalize, struggling when network capacity is overwhelmed by unique inputs. This highlights fundamental limitations of current methods when facing diverse visuals where assumptions break and the learning signal is too weak for generalization.

5.3. Can agents generalize to unseen visual inputs?

Our analysis of out-of-distribution (OOD) generalization capabilities on PPO and SAC agents reveals a stark contrast between in-distribution and generalization performance. We evaluate two levels of generalization difficulty: ‘Easy’ OOD, using augmented versions of training images, and ‘Hard’ OOD, using completely unseen images. Since meaningful generalization evaluation requires agents that first achieve competent in-distribution performance, we primarily focus our OOD analysis on smaller pool sizes (i.e., 1, 5, and 10) where most methods successfully learn the task.

Despite most agents achieving high success rates on their training distribution, they exhibit varying degrees of failure on both generalization challenges. Table 2 illustrates this for the Easy OOD setting across pool sizes 1, 5, and 10, showing a general decline in OOD performance as training pool diversity increases. For Easy OOD evaluation, we applied the augmentations described in Appendix A.3 to training images and tested PPO and SAC agents across 100 episodes

Table 2. Success rate of PPO and SAC agents on Easy OOD across different pool sizes. Higher is better.

Algorithm	Pool 1	Pool 5	Pool 10
PPO	0.49±0.13	0.53±0.14	0.34±0.08
PPO + PT (ID)	0.33±0.09	0.53±0.16	0.27±0.07
PPO + PT (OOD)	0.49±0.12	0.52±0.14	0.34±0.08
SAC	0.45±0.12	0.58±0.12	0.46±0.12
SAC + AE	0.78±0.11	0.64±0.16	0.55±0.12
SAC + VAE	0.64±0.15	0.30±0.08	0.12±0.03
SAC + SPR	0.65±0.13	0.21±0.09	0.07±0.04
SAC + DBC	0.44±0.13	0.34±0.13	0.13±0.04
SAC + CURL	0.76±0.09	0.44±0.10	0.37±0.11
SAC + RAD	0.62±0.15	0.42±0.13	0.30±0.11
SAC + SB	0.89±0.08	0.65±0.12	0.06±0.02

for each augmentation type. We observe that while some SAC variants demonstrate improved robustness to these augmentations compared to base SAC, particularly on pool size 1, this Easy OOD performance generally degrades as the training pool size increases. We found a strong correlation (Pearson $r=-0.81$, $p=2.5e-12$) between success rates on these augmented images (Easy OOD) and agents’ in-distribution sample efficiency, suggesting that representations robust to simple transformations are tied to faster learning on SPGym.

For Hard OOD evaluation, we test PPO, SAC and Dreamer agents on 100 episodes using images independently sampled from the ImageNet-1k validation split, repeating across 5 seeds. Here, agents almost universally fail, achieving near 0% success rates across all methods and training configurations. This stark failure on completely unseen images is a crucial diagnostic finding, exposing fundamental limitations of current end-to-end RL methods when faced with visual inputs that differ substantially from their training distribution. The complete inability to generalize suggests that what agents learn to achieve good task performance is primarily memorization of specific visual patterns rather than learning generalizable representations.

Intuitively, one might expect training on larger, more diverse image pools to improve generalization. However, our results demonstrate that even agents trained on pools of up to 100 images completely fail to transfer skills to novel images in the Hard OOD setting. Furthermore, performance on augmented training images often *decreases* as training pool size increases. This counter-intuitive result may suggest that agents trained on smaller, less diverse pools (which are often more sample efficient) learn representations more attuned to the specific structural invariances of the task, making them robust to simple perturbations but not fundamentally general. The rare non-zero success rates in Hard OOD likely stem from chance encounters with nearly-solved initial states rather than genuine generalization.

These findings offer critical insights for developing future methods. The strong link between Easy OOD robustness and sample efficiency suggests that encouraging invariance to simple transformations is beneficial. However, the universal failure on Hard OOD, and the degradation of Easy OOD with increasing pool diversity, indicates that current methods do not learn sufficiently abstract or disentangled representations. Future research should focus on developing techniques that not only achieve good sample efficiency but also explicitly promote generalization to truly novel visual contexts. This might involve architectures that better separate visual representation learning from policy learning, incorporate stronger inductive biases for visual reasoning, or leverage self-supervised objectives that encourage learning of more fundamental visual features. Simply increasing the diversity of training images, as SPGym allows, is insufficient with current algorithms, highlighting the need for new approaches to bridge this generalization gap. Full OOD performance data for both Easy and Hard settings across all configurations can be found in Appendix C.2.

5.4. Is representation quality linked with performance?

To measure the quality of learned representations, we performed linear probing on frozen encoders from trained PPO and SAC agents, using a single-layer MLP to predict one-hot puzzle states. We find a statistically significant correlation between probe accuracy and sample efficiency (Pearson $r=0.81$, $p=1.1e-13$), indicating that encoders capturing more task-relevant spatial information are strongly linked with faster learning. As image pool size increases, both probe accuracy and task performance systematically degrade, with standard SAC maintaining high accuracy (100% at pool size 1, 97.63% at pool size 5) mirroring its strong efficiency, while less efficient methods like SAC+VAE (78.21% at pool size 5) and SPR (dropping from 94.31% to 75.48% from pool size 5 to 10) show reduced probe performance. These consistent trends across algorithms demonstrate SPGym’s ability to identify learning procedures that develop representations better aligned with the task’s spatial reasoning needs. Full results are provided in Appendix C.1.2.

5.5. Does performance generalize across image sources?

To validate that our findings generalize beyond ImageNet, we evaluated agents on DiffusionDB (Wang et al., 2023), a dataset of procedurally generated images. As shown in Figure 8, performance scaling patterns on DiffusionDB closely mirror those observed on ImageNet across PPO, SAC, and DreamerV3. This consistency across fundamentally different image sources, real photographs versus synthetic generations, demonstrates that visual diversity rather than semantic content drives the representation learning challenge in SPGym. The similarity in degradation patterns as pool size increases indicates that our algorithmic insights reflect

fundamental properties of the tested methods rather than dataset-specific artifacts. Procedurally generated images also offer practical advantages for future research: eliminating storage requirements through on-demand generation, enabling fine-grained control over visual similarity, and providing unlimited training diversity. For a more detailed analysis, see Appendix A.5.

5.6. How does puzzle size affect learning performance?

Another direction is increasing puzzle size. As shown in Table 3, the complexity increase from 3×3 to 4×4 grids significantly impacts learning. On the simpler 3×3 puzzle, PPO solved the puzzle in 1.75M steps, while SAC and DreamerV3 were more efficient, requiring 0.33M and 0.42M steps, respectively. For the 4×4 puzzle, PPO’s sample requirements surged to 24.46M steps, far exceeding the 10M step training budget, while SAC and DreamerV3 still solved the puzzle within budget at 8.14M and 2.26M steps, respectively. This demonstrates that while larger state spaces pose a major challenge by requiring more exploration and more complex visual representations, more sample-efficient algorithms can still scale to harder tasks.

Table 3. Million steps to reach 80% success rate across grid sizes, with pool size 1. Lower is better.

Grid Size	PPO	SAC	DreamerV3
3x3	1.75±0.44	0.33±0.07	0.42±0.06
4x4	24.46±7.58	8.14±3.64	2.26±0.29

5.7. How optimal are the learned solutions?

While our primary focus is on sample efficiency for task completion, we also analyze solution quality by examining the average number of steps agents take to solve puzzles. Our experimental design uses early termination when agents achieve 100% success rate to enable out-of-distribution evaluation before extreme encoder overfitting and to save computational resources. However, this approach may prevent agents from discovering more optimal solutions through continued training. To investigate this trade-off, we trained PPO, SAC, and DreamerV3 on pool size 1 for the full 10M steps without early termination across 5 seeds. Comparing episode lengths between when these agents first achieve 100% success rate (where early termination would occur) and after completing the full training reveals substantial improvements in solution efficiency. For PPO, the first 100 successful episodes average 214.30 ± 16.52 steps, while the last 100 episodes average 31.35 ± 6.59 steps. SAC shows improvement from 64.16 ± 9.81 to 57.27 ± 12.29 steps, and DreamerV3 improves from 126.02 ± 17.25 to 23.48 ± 0.71 steps. Notably, DreamerV3 with continued training approaches the theoretical 22-step average optimal solution (Reinefeld, 1993). This confirms that early termina-

tion, while needed for our experimental objectives, prevents agents from discovering more optimal solutions through continued training.

6. Conclusion

We introduce the Sliding Puzzles Gym (SPGym), a novel benchmark designed to systematically evaluate representation learning in RL algorithms by isolating visual complexity from environment dynamics. Our analysis reveals three fundamental tensions in visual RL that challenge current methods. First, many sophisticated representation learning techniques struggle with SPGym’s unique combination of visual diversity and structured spatial reasoning, with auxiliary objectives often underperforming compared to simpler approaches like data augmentation. Second, we uncover a complex relationship between sample efficiency and generalization, where agents trained on smaller pools often learn representations more attuned to task-specific invariances, making them paradoxically more robust to simple perturbations. Most critically, we expose severe limitations in current end-to-end RL methods through their universal failure on completely novel visual contexts, achieving near-zero success rates despite strong in-distribution performance – revealing that learned representations rely primarily on memorization rather than genuine visual understanding. These findings highlight the need for fundamental algorithmic advances that effectively balance sample efficiency with true visual generalization.

Limitations. We identify two key limitations in this work. First, our objective was to evaluate out-of-the-box performance of all methods with minimal tuning, which means we may not have seen the true peak capabilities of each approach. Second, computational constraints limited us to 5 independent runs per configuration. Given the high stochasticity in sampled image pools, more seeds would provide better statistical robustness for comparing methods.

Future work. SPGym opens several promising research directions for developing representation learning methods that can bridge the generalization gap while maintaining sample efficiency. The benchmark’s controlled nature enables systematic investigation of the trade-offs between training diversity and generalization capability, while its easy integration with external image generation systems allows controlled experiments with varying degrees of visual similarity. Beyond images, SPGym supports incorporating other modalities as puzzle observations, enabling research on leveraging powerful pretrained models as representation encoders. We believe these directions would help develop more robust learning approaches that transfer effectively across different domains and modalities.

Acknowledgements

The authors gratefully acknowledge the valuable insights and constructive discussions provided by Professors Marcos R. O. A. Maximo and Flávio H. T. Vieira.

This work has been partially funded by the project Research and Development of Digital Agents Capable of Planning, Acting, Cooperating and Learning supported by Advanced Knowledge Center in Immersive Technologies (AKCIT), with financial resources from the PPI IoT/Manufatura 4.0 / PPI HardwareBR of the MCTI grant number 057/2023, signed with EMBRAPPI.

Luckeciano C. Melo acknowledges funding from the Air Force Office of Scientific Research (AFOSR) European Office of Aerospace Research & Development (EOARD) under grant number FA8655-21-1-7017.

Impact Statement

This work advances our understanding of how visual representation learning affects reinforcement learning performance and generalization. The insights gained could influence the development of more robust and efficient RL systems for real-world applications. However, several potential impacts warrant consideration:

Positive impacts: By exposing limitations in current methods, our benchmark may help direct research toward more reliable and generalizable RL systems. This could accelerate progress in areas like robotics and autonomous systems where visual understanding is crucial. The benchmark’s controlled nature also promotes more rigorous evaluation practices in RL research.

Negative impacts: As with any machine learning benchmark, there is a risk of overfitting to our specific evaluation criteria rather than addressing fundamental challenges. Additionally, improved visual RL capabilities could enable automation that displaces human workers or enable surveillance applications if misused.

Mitigations: We have open-sourced our benchmark and evaluation protocols to promote transparency and reproducibility. We encourage future work to consider both performance metrics and broader societal implications when building on our findings. Researchers should carefully consider potential dual-use applications and implement appropriate safeguards when deploying systems based on these techniques.

References

- Agostinelli, F., McAleer, S., Shmakov, A., and Baldi, P. Solving the rubik’s cube with deep reinforcement learning and search. *Nature Machine Intelligence*, 1(8):356–363, 2019. doi: 10.1038/s42256-019-0070-z.
- Beattie, C., Leibo, J. Z., Teplyashin, D., Ward, T., Wainwright, M., Küttler, H., Lefrancq, A., Green, S., Valdés, V., Sadik, A., et al. Deepmind lab. *arXiv preprint arXiv:1612.03801*, 2016.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8): 1798–1828, 2013.
- Brockman, G. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Burns, E., Hatem, M., Leighton, M., and Ruml, W. Implementing fast heuristic search code. *Proceedings of the International Symposium on Combinatorial Search*, 3(1): 25–32, 2012.
- Christodoulou, P. Soft actor-critic for discrete action settings. *arXiv preprint arXiv:1910.07207*, 2019.
- Cobbe, K., Hesse, C., Hilton, J., and Schulman, J. Leveraging procedural generation to benchmark reinforcement learning. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pp. 2048–2056. PMLR, 13–18 Jul 2020.
- Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., and Koltun, V. CARLA: An open urban driving simulator. In Levine, S., Vanhoucke, V., and Goldberg, K. (eds.), *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pp. 1–16. PMLR, 13–15 Nov 2017.
- Estermann, B., Lanzendörfer, L. A., Niedermayr, Y., and Wattenhofer, R. Puzzles: A benchmark for neural algorithmic reasoning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- Ghosh, P., Sajjadi, M. S. M., Vergari, A., Black, M. J., and Schölkopf, B. From variational to deterministic autoencoders. In *8th International Conference on Learning Representations (ICLR)*, April 2020.

- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *35th International Conference on Machine Learning (ICML)*, pp. 1861–1870. PMLR, 2018a.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., and Levine, S. Soft actor-critic algorithms and applications. *ArXiv*, abs/1812.05905, 2018b.
- Hafner, D., Pasukonis, J., Ba, J., and Lillicrap, T. Mastering diverse control tasks through world models. *Nature*, 640(8059):647–653, April 2025. ISSN 1476-4687. doi: 10.1038/s41586-025-08744-2.
- Higgins, I., Pal, A., Rusu, A., Matthey, L., Burgess, C., Pritzel, A., Botvinick, M., Blundell, C., and Lerchner, A. Darla: Improving zero-shot transfer in reinforcement learning. In *34th International Conference on Machine Learning (ICML)*, pp. 1480–1490. PMLR, 2017.
- Huang, S., Dossa, R. F. J., Ye, C., Braga, J., Chakraborty, D., Mehta, K., and Araújo, J. G. Cleanrl: High-quality single-file implementations of deep reinforcement learning algorithms. *Journal of Machine Learning Research*, 23(274):1–18, 2022.
- Johnson, W. W. and Story, W. E. Notes on the ‘15’ puzzle. *American Journal of Mathematics*, 2(4):397–404, 1879. doi: 10.2307/2369492.
- Korf, R. E. Depth-first iterative-deepening: An optimal admissible tree search. *Artificial intelligence*, 27(1):97–109, 1985.
- Laskin, M., Lee, K., Stooke, A., Pinto, L., Abbeel, P., and Srinivas, A. Reinforcement learning with augmented data. *Advances in Neural Information Processing Systems*, 33: 19884–19895, 2020a.
- Laskin, M., Srinivas, A., and Abbeel, P. Curl: Contrastive unsupervised representations for reinforcement learning. In *37th International Conference on Machine Learning (ICML)*, pp. 5639–5650. PMLR, 2020b.
- Lee, S. C. and See, T. H. Comparing the hamming and manhattan heuristics in solving the 8—puzzle by a* algorithm. In Mustapha, A. B., Shamsuddin, S., Zuhair Haider Rizvi, S., Asman, S. B., and Jamaian, S. S. (eds.), *Proceedings of the 7th International Conference on the Applications of Science and Mathematics 2021*, pp. 189–195, Singapore, 2022. Springer Nature Singapore. ISBN 978-981-16-8903-1.
- Lesort, T., Díaz-Rodríguez, N., Goudou, J.-F., and Filliat, D. State representation learning for control: An overview. *Neural Networks*, 108:379–392, 2018.
- Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., and Talwalkar, A. Hyperband: A novel bandit-based approach to hyperparameter optimization. *Journal of Machine Learning Research*, 18(185):1–52, 2018.
- Mnih, V., Kavukcuoglu, K., Silver, D., et al. Human-level control through deep reinforcement learning. *Nature*, 2015.
- Moon, S. U. and Cho, Y. Improving the training performance of dqn model on 8-puzzle environment through pre-training. *International Journal of Applied Engineering Research*, 6(1):63–67, 2024.
- Reinefeld, A. Complete solution of the eight-puzzle and the benefit of node ordering in ida*. In *International Joint Conference on Artificial Intelligence*, pp. 248–253. Citeseer, 1993.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C., and Fei-Fei, L. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Schwarzer, M., Anand, A., Goel, R., Hjelm, R. D., Courville, A. C., and Bachman, P. Data-efficient reinforcement learning with self-predictive representations. In *International Conference on Learning Representations*, 2020.
- Schwarzer, M., Rajkumar, N., Noukhovitch, M., Anand, A., Charlin, L., Hjelm, D., Bachman, P., and Courville, A. C. Pretraining representations for data-efficient reinforcement learning. In *Advances in Neural Information Processing Systems*, 2021.
- Stone, A., Ramirez, O., Konolige, K., and Jonschkowski, R. The distracting control suite—a challenging benchmark for reinforcement learning from pixels. *arXiv preprint arXiv:2101.02722*, 2021.
- Stooke, A., Lee, K., Abbeel, P., and Laskin, M. Decoupling representation learning from reinforcement learning. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9870–9879. PMLR, 18–24 Jul 2021.
- Tassa, Y., Doron, Y., Muldal, A., Erez, T., Li, Y., Casas, D. d. L., Budden, D., Abdolmaleki, A., Merel, J., Lefrancq, A., et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018.

- Tomar, M., Mishra, U. A., Zhang, A., and Taylor, M. E. Learning representations for pixel-based control: What matters and why? *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.
- Tomilin, T., Fang, M., Zhang, Y., and Pechenizkiy, M. Coom: A game benchmark for continual reinforcement learning. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 67794–67832. Curran Associates, Inc., 2023.
- Wang, Z. J., Montoya, E., Munechika, D., Yang, H., Hoover, B., and Chau, D. H. DiffusionDB: A large-scale prompt gallery dataset for text-to-image generative models. In Rogers, A., Boyd-Graber, J., and Okazaki, N. (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 893–911, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.51.
- Yarats, D., Kostrikov, I., and Fergus, R. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *9th International Conference on Learning Representations (ICLR)*, 2021a.
- Yarats, D., Zhang, A., Kostrikov, I., Amos, B., Pineau, J., and Fergus, R. Improving sample efficiency in model-free reinforcement learning from images. *Proceedings of the aaai conference on artificial intelligence*, 35(12): 10674–10681, 2021b.
- Zhang, A., McAllister, R. T., Calandra, R., Gal, Y., and Levine, S. Learning invariant representations for reinforcement learning without reconstruction. In *International Conference on Learning Representations*, 2021.

A. Experimental Setup and Configuration

A.1. Model Architectures

We base our implementations on CleanRL’s Atari agents for both PPO and SAC, with minor architectural modifications including additional normalization layers and increased network depth to approximate our SAC implementation to the one used by Tomar et al. (2023). The architectures are detailed below.

For PPO agents:

```
SharedEncoder(
  (encoder): Sequential(
    (0): Conv2d(3, 32, kernel_size=(8, 8), stride=(4, 4))
    (1): ReLU()
    (2): BatchNorm2d(32, eps=1e-05, momentum=0.1, affine=True, stats=True)
    (3): Conv2d(32, 64, kernel_size=(4, 4), stride=(2, 2))
    (4): ReLU()
    (5): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, stats=True)
    (6): Conv2d(64, 64, kernel_size=(3, 3), stride=(1, 1))
    (7): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, stats=True)
    (8): Flatten(start_dim=1, end_dim=-1)
  )
  (projection): Sequential(
    (0): Linear(in_features=3136, out_features=512, bias=True)
    (1): LayerNorm((512,)), eps=1e-05, elementwise_affine=True)
    (2): Tanh()
  )
)
Actor(
  (encoder): SharedEncoder
  (mlp): Linear(in_features=512, out_features=4, bias=True)
)
Critic(
  (encoder): SharedEncoder
  (mlp): Linear(in_features=512, out_features=1, bias=True)
)
```

For SAC agents:

```
SharedEncoder(
  (encoder): Sequential(
    (0): Conv2d(3, 32, kernel_size=(8, 8), stride=(4, 4))
    (1): ReLU()
    (2): BatchNorm2d(32, eps=1e-05, momentum=0.1, affine=True, stats=True)
    (3): Conv2d(32, 64, kernel_size=(4, 4), stride=(2, 2))
    (4): ReLU()
    (5): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, stats=True)
    (6): Conv2d(64, 64, kernel_size=(3, 3), stride=(1, 1))
    (7): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, stats=True)
    (8): Flatten(start_dim=1, end_dim=-1)
  )
)
Actor(
  (encoder): Encoder(
    (shared_encoder): SharedEncoder
    (projection): Sequential(
      (0): Linear(in_features=3136, out_features=512, bias=True)
      (1): LayerNorm((512,)), eps=1e-05, elementwise_affine=True)
      (2): Tanh()
    )
  )
  (mlp): Sequential(
    (0): Linear(in_features=512, out_features=512, bias=True)
    (1): LayerNorm((512,)), eps=1e-05, elementwise_affine=True)
    (2): ReLU()
  )
)
```

```

        (3): Linear(in_features=512, out_features=4, bias=True)
    )
)
Critic(
  (encoder): Encoder(
    (shared_encoder): SharedEncoder
    (projection): Sequential(
      (0): Linear(in_features=3136, out_features=512, bias=True)
      (1): LayerNorm((512,), eps=1e-05, elementwise_affine=True)
      (2): Tanh()
    )
  )
  (mlp): Sequential(
    (0): Linear(in_features=512, out_features=512, bias=True)
    (1): LayerNorm((512,), eps=1e-05, elementwise_affine=True)
    (2): ReLU()
    (3): Linear(in_features=512, out_features=4, bias=True)
  )
  (mlp): Sequential(
    (0): Linear(in_features=512, out_features=512, bias=True)
    (1): LayerNorm((512,), eps=1e-05, elementwise_affine=True)
    (2): ReLU()
    (3): Linear(in_features=512, out_features=4, bias=True)
  )
)
)

```

For the representation learning methods, we maintain the same base architecture and use a consistent MLP structure (2 layers with ReLU activation) for the projector, predictor, transition and reward models. The decoder architecture is as follows:

```

ImageDecoder(
  (decoder): Sequential(
    (0): Linear(in_features=512, out_features=512, bias=True)
    (1): LayerNorm((512,), eps=1e-05, elementwise_affine=True)
    (2): ReLU()
    (3): Linear(in_features=512, out_features=3136, bias=True)
    (4): Unflatten(dim=1, unflattened_size=(3, 7, 7))
    (5): ConvTranspose2d(3, 64, kernel_size=(3, 3), stride=(1, 1))
    (6): ReLU()
    (7): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, stats=True)
    (8): ConvTranspose2d(64, 32, kernel_size=(4, 4), stride=(2, 2))
    (9): ReLU()
    (10): BatchNorm2d(32, eps=1e-05, momentum=0.1, affine=True, stats=True)
    (11): ConvTranspose2d(32, 3, kernel_size=(8, 8), stride=(4, 4))
    (12): Sigmoid()
  )
)

```

For DreamerV3 agents, we use the same base 12M architecture as the one used by [Hafner et al. \(2025\)](#). Appendix A.2 contains the specific hyperparameters used in our experiments.

A.2. Hyperparameters

Table 5, Table 6, and Table 7 list hyperparameters used across all experiments, unless noted otherwise. For DreamerV3, we adopted hyperparameters from ([Hafner et al., 2025](#)), modifying only the decoder loss scale (set to 0) for the version without decoder. Table 8 lists hyperparameters for representation learning methods and components. We use a separate optimizer for the representation learning gradient flow, and we adopt a higher learning rate. When using representation learning methods, we update the target network’s encoder faster, with an EMA τ of 0.025. When using a crop augmentation, we set the image size to 100 and crop it back to 84.

Table 4. Benchmark settings

Parameter	Value
Max steps	10M
Puzzle size	3x3
Action space	discrete
Variation	image
Render size	100x100 (for crop augmentation) 84x84 (otherwise)
Dataset	ImageNet-1k validation split

Table 5. Hyperparameters for PPO

PPO Parameter	Value
Input image size	84x84
Env instances	64
Optimizer	Adam
Learning Rate (LR)	2.5e-4
LR annealing	yes
Adam ϵ	1e-5
Num. steps	16
Num. epochs	4
Batch size	64
Num. minibatches	4
γ	0.99
GAE λ	0.95
Advantage normalization	yes
Clip coef.	0.1
Clip value loss	yes
Value function coef.	0.5

Table 6. Hyperparameters for SAC

SAC Parameter	Value
Input image size	84x84
Env instances	64
Optimizer	Adam
Learning rate	3e-4
Replay buffer capacity	3e5
Batch size	4096
Warmup steps	2e4
γ	0.99
Policy update frequency	2
Fixed α temperature	0.05
Target network update frequency	1
Target Q functions EMA τ	0.005
Target encoder EMA τ	0.005 (standard) 0.025 (otherwise)

Table 7. Hyperparameters for DreamerV3

DreamerV3 Parameter	Value
Input image size	80x80
Env instances	16
Model size	12M
RSSM deterministic size	2048
RSSM hidden size	256
RSSM classes	16
Network depth	16
Network units	256
Replay buffer capacity	5e5
Replay ratio	32
Action repeat	1
Learning rate	4e-5
Batch size	16
Batch length	64
Imagination horizon	15
Discount horizon	333
Decoder loss scale	1 (standard) 0 (no decoder)

Table 8. Hyperparameters for Representation Learning Methods and Components

Method	Hyperparameter	Value
General	Learning rate	1e-3
	Loss coefficient	1.0
Transition and reward models	Min sigma	1e-4
	Max sigma	10
	Probabilistic	no
SAC-AE	Latent space decay weight	1e-6
	Decoder decay weight	1e-7
SAC-VAE	Variational KL weight β	1e-7
CURL	Temperature	0.1
	Positive samples	temporal/augmented
	Augmentations	crop
RAD	Augmentations	crop channel_shuffle color_jitter
	Horizon H	5
SPR	Augmentations	crop

A.3. Augmentation Strategies

We evaluated several image augmentations in our preliminary experiments, as described in Appendix A.4.2. These augmentations are illustrated in Figure 5, and are as follows:

- **No augmentation:** The image is fed as is to the agent.
- **Crop:** Randomly crops a portion of the image and resizes it back to the original dimensions. This helps learn translation invariance by forcing the agent to recognize patterns regardless of their position in the frame.

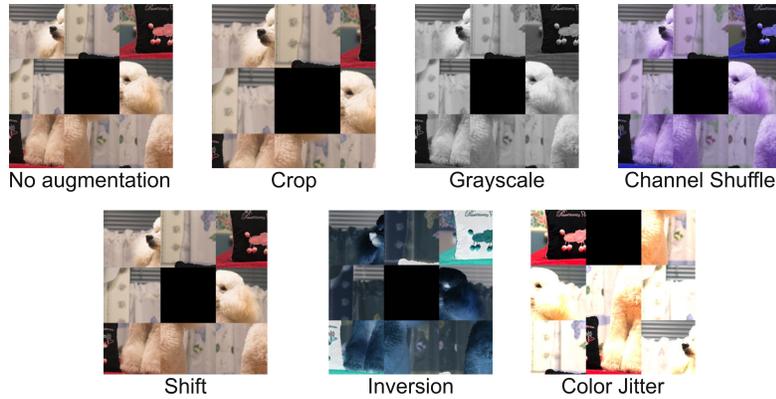


Figure 5. **Example of image augmentations.** The top left image shows the original observation, and the subsequent images show the observation after each augmentation procedure is independently applied.

- **Grayscale:** Converts the RGB image to grayscale by averaging across color channels. This reduces the visual complexity and helps the agent focus on structural features rather than color information.
- **Channel Shuffle:** Randomly permutes the RGB color channels. This encourages the agent to be invariant to color transformations while preserving the image structure.
- **Shift:** Translates the image by a small random amount in both horizontal and vertical directions. Similar to crop, this promotes translation invariance in the learned representations.
- **Inversion:** Inverts the pixel values by subtracting them from the maximum possible value (255 for 8-bit images). This teaches the agent to recognize patterns independent of absolute intensity values.
- **Color Jitter:** Applies random color variations to the image, including brightness, contrast, saturation, and hue. This helps the agent to be invariant to color transformations while preserving the image structure.

After extensive experimentation, we found that the combination of grayscale and channel shuffle consistently produced the best results. This combination effectively reduces visual complexity while maintaining important structural information. We adopted this augmentation pair as the standard for all our agents that use augmentation techniques.

A.4. Preliminary Experiments and Design Rationale

A.4.1. HYPERPARAMETER SELECTION PROCESS

For SAC agents, the entropy coefficient α significantly impacts performance (Haarnoja et al., 2018a). While Haarnoja et al. (2018b) proposed automatic tuning (autotune) based on policy entropy, we found this approach ineffective for SPGym, even with various learning rates (LRs). Through systematic Hyperband (Li et al., 2018) sweeps over pools of size 1, we identified $\alpha = 0.05$ as optimal (Figure 6). This value provides a good balance between exploration and exploitation, allowing the agent to efficiently learn the puzzle mechanics while maintaining enough randomness to discover new solutions.

PPO and DreamerV3 agents proved more robust to hyperparameter choices, performing well with their default configurations. This robustness is particularly valuable in our benchmark setting, as it suggests these algorithms can

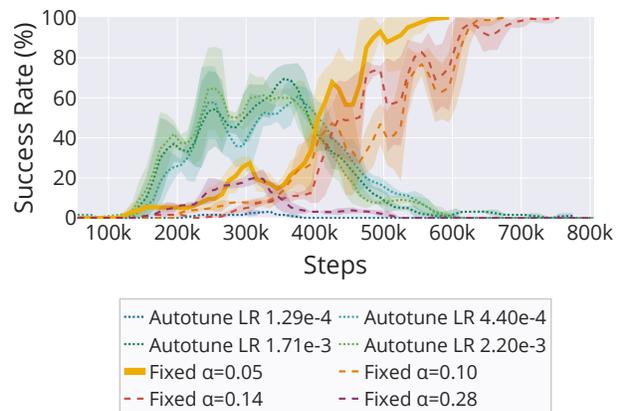


Figure 6. **Performance comparison of different α values for SAC agents on pool size 1.** Fixed $\alpha = 0.05$ outperforms automatic tuning approaches across different learning rates.

adapt to new tasks without extensive tuning. Complete hyperparameter settings are provided in Appendix A.2.

A.4.2. DATA AUGMENTATION ANALYSIS AND CHOICE

Building on insights from Laskin et al. (2020a) that environment-specific invariances influence optimal augmentation strategies, we systematically evaluated augmentation pipelines for RAD, CURL, and SPR in SPGym. Our analysis focused on sample efficiency to maintain consistency with our core experimental objectives. We tested the five augmentation techniques detailed on Appendix A.3 on pools of 5 images, which offered a good balance between task complexity and convergence speed. For SPR, specifically, we also experimented with shift + color jitter, as suggested by Schwarzer et al. (2020). Across all algorithms, our experiments consistently converged to a simple two-stage augmentation process: probabilistic grayscale conversion (with 20% chance) followed by channel shuffling (Figure 7). This pipeline’s effectiveness likely stems from its ability to simultaneously reduce visual complexity through grayscale conversion while introducing beneficial stochasticity via channel shuffling, with both transformations preserving critical structural information while preventing overfitting to specific color patterns. We adopted this augmentation combination for all subsequent experiments. While our current evaluation focused on sample efficiency, investigating how different augmentation strategies affect generalization remains an important direction for future work.

A.5. Dataset Analysis and Choice

Our choice of ImageNet-1k as the primary dataset for SPGym was motivated by several key considerations. First, ImageNet provides a diverse set of real-world images that challenge agents to learn generalizable visual representations. We hypothesized that real-world images would provide unique insights related to representation learning that should be applicable to other domains beyond the puzzle proposed in SPGym, as they contain the complex visual patterns and structures that agents encounter in practical applications. However, as shown in Figure 8 and in comparison to Figure 4, the performance scaling patterns we observe on ImageNet closely mirror those on DiffusionDB, suggesting that our findings are not specific to a particular dataset but rather reflect fundamental properties of the algorithms being tested.

The similarity in scaling behavior between ImageNet and DiffusionDB is particularly noteworthy because these datasets differ substantially in their composition and generation process. While ImageNet consists of real photographs, DiffusionDB contains synthetic images generated by text-to-image models. The consistent performance patterns across these datasets suggest that our results capture fundamental algorithmic behaviors rather than dataset-specific artifacts.

Our demonstration with DiffusionDB reveals promising directions for future work with procedurally generated datasets. As shown in our analysis, agents perform similarly on ImageNet and DiffusionDB, suggesting that visual diversity rather than semantic content drives difficulty. Procedurally generated images offer several compelling advantages worth further investigation: they eliminate the need for large image storage by generating unique images on demand for each episode, enable fine-grained control over the generalization challenge by gradually increasing visual differences between generated images, and provide virtually unlimited training data diversity. These capabilities could enable more systematic studies of

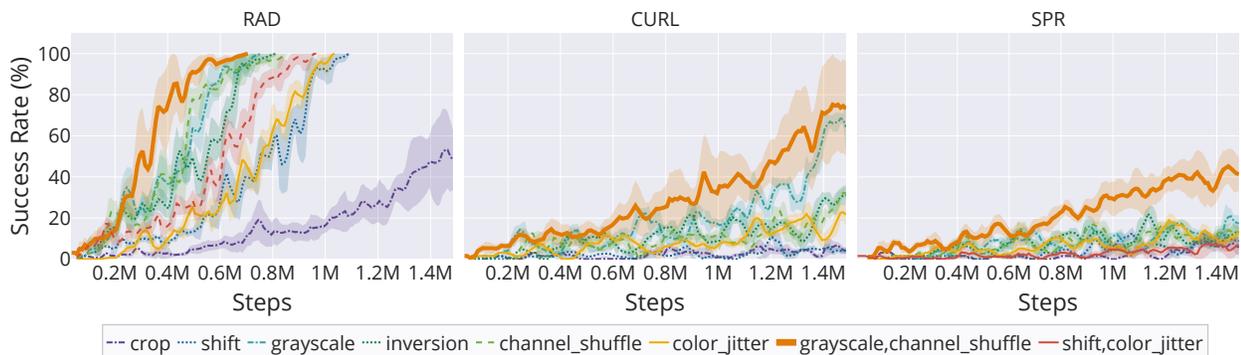


Figure 7. Comparative analysis of data augmentation strategies. Results show performance of SAC with RAD, CURL and SPR on 5-image pools. Grayscale conversion and channel shuffling emerge as the most effective combination, significantly outperforming other augmentation strategies.

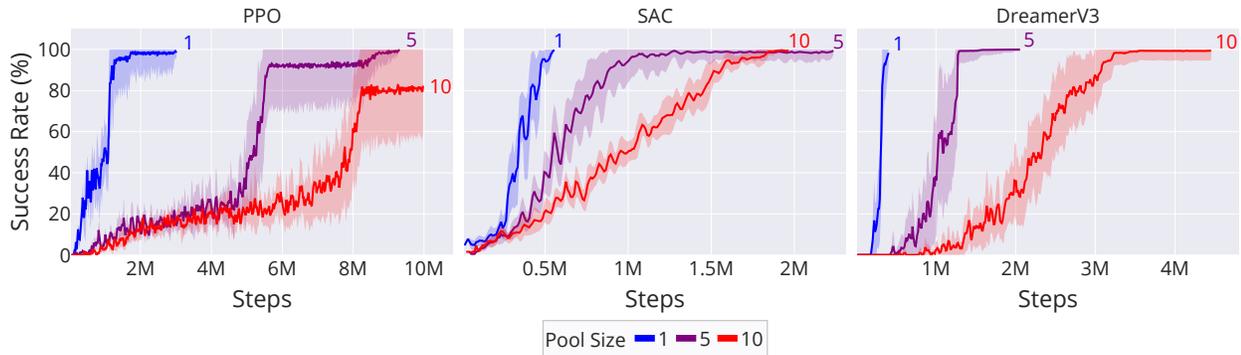


Figure 8. **Performance scaling with DiffusionDB images.** Success rates for PPO, SAC, and DreamerV3 agents across different pool sizes (1, 5, and 10) using DiffusionDB images. The performance patterns closely mirror those seen with ImageNet (Figure 4). Shaded regions represent 95% confidence intervals across 5 independent seeds.

visual generalization in reinforcement learning.

This cross-dataset consistency strengthens our confidence in the generalizability of our findings. It indicates that the relative performance of different representation learning methods is driven more by their core assumptions and architectural choices than by the specific characteristics of the training data. This is particularly important for our goal of understanding how different approaches handle increasing visual diversity, as it suggests our conclusions may extend to other domains beyond the specific datasets used in our experiments.

A.6. Hardware Setup and Runtime

Our hardware setup consists of an AMD Ryzen 7 3700X CPU, an NVIDIA RTX 3090 GPU, 64GB of RAM, and 128GB of swap space. Using this configuration, DreamerV3 experiments take approximately 20 hours per run, primarily because of the heavy use of swap space for replay buffers, which must store hundreds of thousands of images in memory. SAC takes between 2 (e.g. standard, RAD) to 11 hours (e.g. SPR, SB) depending on the representation learning components used. For SAC, the longer runtimes are due to the number of sequential inference steps required to train the agent and auxiliary networks. PPO experiments are significantly faster, with the longest runs completing in about 1 hour and 30 minutes.

B. Algorithms and Variations

We describe the algorithms and their variations used in this work in detail. Our experiments employ three main algorithms – PPO, SAC, and DreamerV3 – each with different representation learning approaches. We focus particularly on how these methods process and learn from visual observations, as this is crucial for performance in our benchmark.

B.1. Pretraining and PPO

Drawing inspiration from previous work on pretraining methods in RL (Higgins et al., 2017; Stooke et al., 2021; Schwarzer et al., 2021), we implement a pretraining approach for PPO that focuses on learning task-relevant visual representations. The pretraining process involves training a PPO agent to completion on a single environment instance, then extracting its CNN weights. These pretrained weights are then used to initialize new PPO agents, while all other network components (policy and value networks) start from random initialization. We evaluate two scenarios: in-distribution (ID), where new agents are trained on the same pool of images used during pretraining, and out-of-distribution (OOD), where a different image pool is sampled. The ID setting represents an upper bound on what pretraining can achieve with perfect visual alignment, while the OOD setting reflects the more realistic scenario of deploying pretrained encoders on novel visual inputs, similar to how general-purpose pretrained models would be used in practice.

B.2. Data Augmentation and RAD

Reinforcement Learning from Augmented Data (RAD) (Laskin et al., 2020a) represents a simple yet effective approach to visual representation learning in RL. The key insight is that applying data augmentation to observations can improve sample efficiency by exposing the agent to transformed versions of experienced states. This creates an implicit regularization effect that helps learn more robust representations.

In our implementation, we combine RAD with SAC and apply data augmentation consistently during both policy updates and value function learning. These augmentations are applied to observations sampled from the replay buffer before being processed by the encoder network. The augmentation pipeline consists of two key transformations identified through our preliminary experiments (Appendix A.4.2): grayscale conversion and channel shuffling. Following the SPR approach (Schwarzer et al., 2020), we apply augmentations independently to each transition after sampling batches from the replay buffer, meaning that samples from the same episode may be augmented differently. We detail how each augmentation procedure is implemented in Appendix A.3.

B.3. Contrastive Learning and CURL

Contrastive learning methods learn representations by maximizing similarity between different views of the same observation while minimizing similarity to other observations. In the context of RL, CURL (Laskin et al., 2020b) applies this principle by using data augmentation to create positive pairs, enabling agents to learn invariant representations. The method achieves this by applying random crops to observations and treating differently augmented views of the same observation as positive pairs.

The contrastive loss for a positive pair of observations (x, x^+) and a set of negative examples $\{x_i^-\}$ is formulated as:

$$\mathcal{L}_{CURL} = -\log \frac{\exp(f_\theta(x)^T f_\theta(x^+)/\alpha)}{\exp(f_\theta(x)^T f_\theta(x^+)/\alpha) + \sum_i \exp(f_\theta(x)^T f_\theta(x_i^-)/\alpha)}, \quad (3)$$

where α is a temperature parameter and f_θ is the encoder function.

In our implementation, we combine CURL with SAC and use the same augmentation strategy identified in Appendix A.3 (grayscale conversion and channel shuffling) rather than the random crops from the original CURL paper. The encoder is trained jointly with the RL objective, allowing the representations to adapt to both the contrastive learning task and the control problem. Negative examples are drawn from other observations within the same batch, providing a computationally efficient way to obtain contrastive pairs without requiring additional memory storage.

B.4. State Metrics and DBC

Deep Bisimulation for Control (DBC) (Zhang et al., 2021) takes a different approach to representation learning by focusing on behavioral similarity between states rather than visual similarity. The key idea is to learn an encoder that maps states to a representation space where distances reflect how similarly states behave in terms of rewards and transitions, rather than how visually similar they appear.

Given pairs of observations (x_i, x_j) , DBC trains an encoder f_θ to minimize:

$$J(\phi) = \left(\|\hat{z}_i - \hat{z}_j\|_1 - |r_i - r_j| - \gamma W_2(\hat{\mathcal{P}}(\cdot|\bar{z}_i, a_i), \hat{\mathcal{P}}(\cdot|\bar{z}_j, a_j)) \right)^2, \quad (4)$$

where $\hat{z}_i = f_\theta(x_i)$ represents the encoded state, $\bar{z}_i = sg(f_\theta(x_i))$ is the stop-gradient version of the encoding, and $\hat{\mathcal{P}}$ is a probabilistic transition model that predicts the next state distribution. The W_2 term represents the 2-Wasserstein distance between predicted transition distributions, which for Gaussian distributions has a closed-form solution (Zhang et al., 2021).

In our implementation, we combine DBC with SAC, jointly training the encoder with both the bisimulation objective and the RL objective. The transition model operates in latent space, predicting Gaussian distributions over next states. This approach helps the agent learn representations that capture behaviorally meaningful features while ignoring visual distractors that don't affect the game dynamics. Unlike methods that rely on data augmentation or reconstruction, DBC's focus on behavioral similarity makes it particularly suited for environments where visually different states might require similar actions.

B.5. Reconstruction-Based Methods and SAC-AE/VAE

Reconstruction-based methods learn representations by training an encoder-decoder architecture to compress and reconstruct observations. We evaluate two variants combined with SAC: SAC-AE using a deterministic autoencoder and SAC-VAE using a variational autoencoder (Yarats et al., 2021b).

For SAC-AE, given an observation x from the replay buffer, we train an encoder f_θ and decoder g_ϕ to minimize:

$$\mathcal{L}_{RAE} = \mathbb{E}_{x \sim \mathcal{D}} [\|x - g_\phi(f_\theta(x))\|^2 + \lambda_z \|f_\theta(x)\|_2^2 + \lambda_\phi \|\phi\|_2^2], \quad (5)$$

where λ_z and λ_ϕ are regularization coefficients that help prevent representation collapse and overfitting respectively (Ghosh et al., 2020).

For SAC-VAE, we replace the deterministic encoder with a probabilistic encoder q_ψ that outputs a distribution over latent states. The training objective becomes:

$$\mathcal{L}_{VAE} = \mathbb{E}_{q_\psi(\hat{z}|x)} [\log g_\phi(x|\hat{z})] - \beta D_{KL}(q_\psi(\hat{z}|x) \|\mathcal{N}(0, 1)), \quad (6)$$

where β controls the trade-off between reconstruction quality and latent space regularization.

In both variants, we train the encoder jointly with the SAC objective, allowing the representations to adapt to both reconstruction and control tasks. The encoded states are used as inputs to the policy and value networks. Unlike methods that rely on data augmentation or behavioral similarity, these approaches learn representations by explicitly modeling the visual structure of observations through reconstruction.

B.6. World Models and DreamerV3

World models learn to predict future states and outcomes by learning a compact latent representation of the environment. DreamerV3 (Hafner et al., 2025) represents the state-of-the-art in world model-based reinforcement learning, employing an online encoder f_θ that maps observations x_t into latent states \hat{z}_t . A recurrent dynamics model h_ω operates in this latent space to predict future states conditioned on actions, while a reward predictor estimates immediate rewards.

The model is trained using multiple objectives that create a multi-task learning pressure. The encoder and a corresponding decoder are trained to reconstruct observations, ensuring the latent space captures relevant visual features. The dynamics model is trained to predict future latent states that lead to accurate reconstructions of future observations. This temporal consistency objective forces the representations to be predictive of future states while supporting reconstruction and control.

DreamerV3 introduces several innovations for stable representation learning, including KL balancing to maintain informative latent states and symmetric cross-entropy loss for better gradients. Unlike methods focused solely on visual similarity, world models must learn representations that serve multiple purposes – capturing visual features, encoding dynamics, and providing a suitable space for policy learning. We refer readers to Hafner et al. (2025) for implementation details.

B.7. Temporal Consistency Methods

Several methods leverage temporal consistency in the environment to learn better representations. These approaches are based on the principle that a good representation should not only capture the current state but also be predictive of future states and outcomes.

B.7.1. SELF-PREDICTIVE REPRESENTATIONS (SPR)

SPR (Schwarzer et al., 2020) represents a non-contrastive approach that learns by predicting future latent states. Given a sequence of states and actions $(x_{t:t+K}, a_{t:t+K})$ from the replay buffer, where K is the prediction horizon, SPR employs:

- An online encoder f_θ that maps observations to latent states: $\hat{z}_t = f_\theta(x_t)$
- A target encoder $f_{\theta'}$ providing stable training targets, updated via exponential moving average
- An action-conditioned transition model h_ω that predicts future latent states: $\hat{z}_{t+k+1} = h_\omega(\hat{z}_{t+k}, a_{t+k})$
- Projection networks $p_\xi, p_{\xi'}$ and prediction head w_ζ that transform representations for the prediction task

The model generates predictions $\hat{y}_{t+k} = w_{\zeta}(p_{\xi}(\hat{z}_{t+k}))$ and compares them to target projections $\tilde{y}_{t+k} = p_{\xi'}(\tilde{z}_{t+k})$ using a cosine similarity loss:

$$\mathcal{L}_{\text{SPR}} = - \sum_{k=1}^K \left(\frac{\tilde{y}_{t+k}}{\|\tilde{y}_{t+k}\|_2} \right)^T \left(\frac{\hat{y}_{t+k}}{\|\hat{y}_{t+k}\|_2} \right) \quad (7)$$

B.7.2. SIMPLE BASELINE METHOD

Tomar et al. (2023) take a minimalist approach to temporal consistency by combining two predictive objectives: reward prediction and transition prediction. We refer to this approach as the Simple Baseline (SB) method. While originally intended as a baseline, it demonstrates the effectiveness of basic temporal prediction for representation learning.

The method augments standard RL algorithms with two predictive components in latent space:

- A transition model h_{ω} that predicts the next encoded state
- A reward predictor h_{reward} that estimates immediate rewards

The transition loss \mathcal{L}_{dyn} measures the mean squared error between predicted and actual next encoded states, while the reward loss $\mathcal{L}_{\text{reward}}$ measures the error in reward predictions. These predictive losses are combined with the standard RL objective:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{RL}} + \mathcal{L}_{\text{reward}} + \mathcal{L}_{\text{dyn}} \quad (8)$$

In our implementations, we combine both SPR and SB with SAC. For SPR, we apply the augmentation strategy identified in Appendix A.3 before encoding observations. Both methods operate entirely in latent space, avoiding the computational cost of pixel-space reconstruction while leveraging temporal structure to learn meaningful representations.

C. Supplementary Analyses and Results

C.1. Representation Learning Analysis

We further evaluate SPGym’s representation learning assessment capabilities through two analyses: (1) comparing raw pixel vs. ground-truth state learning, and (2) linear probes of learned representations.

C.1.1. STATE-BASED VS. IMAGE-BASED OBSERVATIONS

To establish a baseline and understand the impact of the visual representation challenge, we trained PPO, SAC, and DreamerV3 agents using SPGym’s one-hot encoding variation. These one-hot vectors represent the ground-truth puzzle state, identical to the targets used for our linear probes (see Appendix C.1.2). We compared their sample efficiency (steps to 80% success, averaged over 5 seeds) against their image-based counterparts. For PPO and SAC agents processing one-hot vectors, CNN encoders were replaced with 2-layer MLPs. DreamerV3 utilized its default non-image encoder, a 3-layer MLP. Hyperparameters were kept consistent with image-based experiments, without specific tuning for the one-hot setting.

Table 9. **Steps to 80% success on one-hot vs. image-based observations.** Lower is better. ‘-’ indicates experiments were not run for that specific configuration due to computational constraints.

Algorithm	Grid Size	One-hot	Image (Pool 1)	Image (Pool 5)
PPO	3x3	661.69k±81.44k	1.75M±444.81k	7.80M±1.08M
	4x4	12.29M±467.84k	24.46M±7.58M	-
SAC	3x3	672.51k±63.10k	334.26k±67.47k	907.21k±116.20k
	4x4	5.09M±463.14k	8.14M±3.64M	-
DreamerV3	3x3	834.86k±61.10k	417.09k±55.03k	1.23M±199.49k
	4x4	3.68M±436.97k	2.26M±287.23k	5.81M ± 2.17M

The results in Table 9 offer several insights. For PPO (both grid sizes) and SAC (4x4 grid), learning directly from ground-truth one-hot states is more sample efficient than learning from images. This is expected, as the one-hot encoding removes

the burden of representation learning from pixels. The instances where image-based agents (SAC and DreamerV3 on 3x3 grid, pool 1) converged faster than their one-hot counterparts might be attributed to differences in network architectures (MLP vs. CNN/Transformer backbones) and the absence of specific tuning for the one-hot setting.

Crucially, the one-hot encoding setting presents a fixed, minimal representation learning challenge. In contrast, SPGym’s image-based variations allow for a systematic scaling of the visual diversity challenge by increasing the image pool size (e.g., Pool 1 vs. Pool 5 vs. Pool 10, etc., see Appendix C.2). Across all agents and grid sizes, increasing visual diversity from pool size 1 to pool size 5 (and beyond) consistently increases sample complexity. This demonstrates SPGym’s ability to isolate and stress the visual representation learning component, as the underlying task dynamics remain constant. This controlled evaluation reveals limitations in how effectively different RL agents learn representations under scalable visual diversity, insights not apparent from the one-hot setting alone. While perfect disentanglement of representation learning from policy learning is challenging in end-to-end training, SPGym provides a valuable framework for structured, comparative evaluation of visual representation learning capabilities in RL.

C.1.2. LINEAR PROBE ANALYSIS

To directly assess the quality of learned visual representations, we conducted linear probing experiments. For each trained PPO and SAC agent, we froze its encoder and trained a single-layer MLP classifier on top of the features extracted by it until convergence. The classifier’s task was to predict the one-hot encoded ground-truth puzzle state corresponding to the input image. This setup allows us to quantify how much task-relevant spatial information is captured by the agent’s encoder. High probe accuracy indicates the encoder has learned features that are linearly separable with respect to the underlying game state.

Our analysis reveals several key insights. First, we find a strong, statistically significant correlation between linear probe accuracy and agent sample efficiency (Pearson $r=-0.81$, $p=1.1e-13$), with higher probe accuracy being highly predictive of fewer environment steps required to reach 80% success. This suggests that agents whose encoders capture more task-relevant spatial information tend to learn the task more efficiently.

Second, examining probe accuracies across pool sizes and algorithms further clarifies this relationship. As image pool size increases, both probe accuracy and task performance systematically degrade, isolating the effect of visual diversity on representation learning. For example, standard SAC maintains high probe accuracy (100% at pool size 1, 97.63% at pool size 5), mirroring its strong sample efficiency. In contrast, methods with lower sample efficiency show reduced probe performance: SAC+VAE achieves only 78.21% probe accuracy at pool size 5, while SAC+RAD reaches 98.66%. Other representation learning methods that underperform standard SAC, such as SPR and DBC, also exhibit declining probe accuracy as pool size increases (e.g., SPR drops from 94.31% at pool size 5 to 75.48% at pool size 10).

These trends indicate that different algorithms develop representations with varying alignment to the spatial reasoning demands of the task. The consistent link between probe accuracy and sample efficiency across diverse methods suggests that SPGym can help identify which learning procedures lead to representations that better support task performance.

Full linear probe accuracy data for all agents and pool sizes are presented in Table 10 below. Comprehensive performance metrics, including the sample efficiency data used in the correlation analysis, can be found in Appendix C.2.

C.2. Detailed Performance Analysis

This section provides comprehensive analysis of algorithmic performance in SPGym, examining how different representation learning approaches handle increasing visual diversity. We present detailed learning curves showing the training dynamics of each method, along with quantitative performance metrics across varying pool sizes.

The analysis includes detailed algorithmic variant analysis explaining why certain methods succeed or fail, and comprehensive performance curves for PPO (Figure 10), SAC (Figure 11), and DreamerV3 (Figure 12). Tables 11 to 13 provide quantitative metrics per pool size (Pool) including sample efficiency (Steps), and episode length (Length), across algorithms and variants. Tables 11 and 12 also include in-distribution and out-of-distribution success rates (ID Success, OOD Easy, OOD Hard), which were not computed for DreamerV3 due to complexities in the original codebase. The rare non-zero success rates in Hard OOD likely stem from chance encounters with nearly-solved initial states rather than genuine generalization.

Table 10. **Linear probe accuracy (%) for each agent and pool size.** Accuracy of the linear probe indicates linearly separable features capturing task-relevant spatial information. There is a statistically significant correlation between probe accuracy and agent sample efficiency (steps to 80% success).

Agent	Pool Size	Steps (M)	Linear Probe Accuracy (%)
PPO	1	1.75±0.44	99.81±0.14
	5	7.80±1.08	96.42±1.06
	10	9.73±0.36	87.84±1.40
	20	10.00±0.00	66.97±5.28
PPO+PT(ID)	1	0.95±0.21	99.81±0.12
	5	5.55±1.22	96.83±0.61
	10	9.17±1.10	89.54±0.95
PPO+PT(OOD)	1	1.34±0.42	99.59±0.33
	5	7.03±1.07	95.68±0.77
	10	9.70±0.41	88.90±1.11
SAC	1	0.33±0.07	100.00±0.00
	5	0.91±0.12	97.63±0.68
	10	1.65±0.31	93.34±0.48
	20	4.52±1.43	80.74±6.31
	30	9.23±0.96	66.69±8.41
	50	10.00±0.00	55.52±0.09
SAC+RAD	1	0.24±0.03	99.99±0.01
	5	0.42±0.06	98.66±0.20
	10	0.82±0.18	89.74±0.73
SAC+CURL	1	0.46±0.10	99.98±0.03
	5	1.56±0.31	97.14±0.17
	10	5.24±1.92	89.47±1.23
SAC+SPR	1	2.09±0.81	99.99±0.01
	5	3.68±1.68	94.31±0.24
	10	10.00±0.00	75.48±1.82
SAC+DBC	1	0.44±0.04	100.00±0.00
	5	0.99±0.25	94.26±1.19
	10	10.00±0.00	76.59±5.82
SAC+AE	1	0.42±0.09	100.00±0.00
	5	1.04±0.24	95.52±4.56
	10	2.03±0.38	88.66±1.88
SAC+VAE	1	1.13±0.14	99.66±0.06
	5	5.30±0.68	78.21±2.35
	10	10.00±0.00	64.76±0.11
SAC+SB	1	0.98±0.88	99.90±0.03
	5	2.08±0.30	96.69±1.08
	10	10.00±0.00	81.93±6.06

C.2.1. ANALYSIS OF ALGORITHMIC VARIANTS

Our preliminary analysis of algorithmic variants suggests how method assumptions may influence effectiveness in SPGym. PPO with in-distribution pretraining (PT (ID)) significantly boosts sample efficiency across all pool sizes, though this advantage diminishes at pool size 10. Out-of-distribution pretraining (PT (OOD)) offers more modest gains that also decrease with larger pools, suggesting limited transfer from general-purpose pretrained encoders. For SAC, data augmentation via RAD consistently improves efficiency across all pool sizes, with particularly pronounced benefits for larger pools. Conversely, many sophisticated auxiliary methods struggle: CURL, SPR, and VAE variants consistently require more samples than standard SAC, with particularly poor performance on larger pools. DBC and AE generally underperform or offer marginal improvements. DreamerV3 demonstrates particularly strong performance, consistently outperforming both PPO and SAC variants across all pool sizes with remarkably stable performance. The variant without decoder gradients shows reduced performance, highlighting the importance of the reconstruction objective and suggesting that learning a predictive environment model provides an effective foundation for handling visual diversity. We now provide more detailed hypotheses for these behaviors.

Pretraining. As shown in Figure 10, pretraining provides clear benefits for PPO, especially with larger pools. While in-distribution pretraining provides strong gains, almost matching the performance of PPO with one-hot-based observations (see Section C.1.1), this represents an optimistic upper bound since real-world scenarios rarely permit task-specific pretraining. Interestingly, out-of-distribution pretraining also shows benefits compared to random initialization (90% vs 86% success at pool size 5), suggesting some transfer of useful visual features from the pretrained encoder. This indicates that even general-purpose pretrained encoders can provide a helpful initialization for RL tasks, though not matching the performance of task-specific pretraining.

Data augmentation. RAD succeeds by enforcing spatial invariances through grayscale+channel shuffling, preserving structural relationships critical for puzzle solving while adding beneficial stochasticity. Its weak assumptions make it robust across diversity levels, maintaining strong and consistent performance across pool sizes through this simple augmentation-based approach.

Contrastive learning. CURL underperforms as instance discrimination may prioritize whole-image features over tile-level details needed for puzzle solving. This suggests contrastive learning’s focus on global image similarity may not align well with the local spatial reasoning required for puzzle solving.

State similarity learning. DBC fails possibly because its core assumption – that states with similar dynamics should have similar representations – breaks down in two ways: identical puzzle states appear radically distinct between episodes with different sampled images, while different states can share visual patterns due to being from the same episode or having the same base image.

Temporal consistency. While the environment’s underlying dynamics are deterministic, temporal consistency methods such as SPR and SB face three key challenges: (1) The visual manifestation of state transitions varies dramatically between episodes due to different base images, forcing the encoder to learn position-invariant representations that capture tile relationships rather than visual content – a difficult disentanglement problem. (2) The assumption of smooth latent space transitions is violated by the discrete nature of tile movements, where single actions induce significant changes in both visual appearance and puzzle state. (3) Most crucially, these methods must simultaneously learn two competing objectives: temporal predictability in latent space (for transition modeling) and visual discriminability (for representation learning). This creates a conflict where features useful for predicting latent transitions (tile positions) are obscured by visually salient but dynamically irrelevant image content. SPR’s prediction horizon mechanism exacerbates this by compounding representation errors through multiple latent transition steps. Similarly, SB’s transition/reward prediction suffers because the latent space conflates visual features with positional information – while rewards depend solely on tile positions, the visual diversity in observations provides no direct positional cues. The same absolute position (e.g., top-left corner) shows completely different visual content each episode, making position-aware latent representations particularly difficult to learn.

Reconstruction-based learning. The success of DreamerV3’s decoder highlights the value of the discrete reconstruction loss in learning useful representations for SPGym. In contrast, simple autoencoders (AE) offer little benefit for SAC, and variational autoencoders (VAE) hurt performance possibly because their continuous latent space assumptions conflict with SPGym’s discrete state transitions.

These findings align with observations from [Tomar et al. \(2023\)](#), who noted that many representation learning methods underperform or fail completely when tested outside their original domain. We note that our evaluation focused on using each algorithm’s suggested hyperparameters for visual RL with discrete actions, aiming to assess their out-of-the-box performance. The results may not reflect the best possible performance achievable through extensive hyperparameter tuning.

Table 11. Performance metrics for PPO agents across image pool sizes. ‘Steps’ indicates the number of environment steps (in millions) required to reach 80% success rate during training. ‘ID Success’ shows the final success rate at the end of 10M training steps when evaluating on the same pool used during training. ‘OOD Easy’ shows the success rate when evaluated on 100 unseen images from the ‘Easy’ distribution. ‘OOD Hard’ shows the success rate when evaluated on 100 unseen images from the ‘Hard’ distribution. ‘Length’ indicates the average number of steps required to solve an instance of the puzzle at the last 100 episodes.

Algorithm	Variation	Pool	Steps (M)	ID Success (%)	OOD Easy (%)	OOD Hard (%)	Length
PPO	Standard	1	1.75±0.44	100±0	0.49±0.13	0.0±0.0	72.9±4.4
		5	7.80±1.08	86±16	0.53±0.14	0.3±0.3	212.0±19.0
		10	9.73±0.36	47±18	0.34±0.08	0.6±0.4	599.6±25.6
		20	10.00±0.00	12±7	0.12±0.03	0.0±0.0	903.4±23.9
	PT (ID)	1	0.95±0.21	100±0	0.33±0.09	0.1±0.1	76.2±4.8
		5	5.55±1.22	100±0	0.53±0.16	0.3±0.4	83.2±7.2
		10	9.17±1.10	38±20	0.27±0.07	0.9±0.7	658.4±25.3
	PT (OOD)	1	1.34±0.42	100±0	0.49±0.12	0.2±0.3	72.7±4.0
		5	7.03±1.07	90±16	0.52±0.14	0.3±0.3	156.0±16.0
		10	9.70±0.41	46±19	0.34±0.08	0.8±1.2	572.0±26.0

Table 12. Performance metrics for SAC agents across image pool sizes. ‘Steps’ indicates the number of environment steps (in millions) required to reach 80% success rate during training. ‘ID Success’ shows the final success rate at the end of 10M training steps when evaluating on the same pool used during training. ‘OOD Easy’ shows the success rate when evaluated on 100 unseen images from the ‘Easy’ distribution. ‘OOD Hard’ shows the success rate when evaluated on 100 unseen images from the ‘Hard’ distribution. ‘Length’ indicates the average number of steps required to solve an instance of the puzzle at the last 100 episodes.

Algorithm	Variation	Pool	Steps (M)	ID Success (%)	OOD Easy (%)	OOD Hard (%)	Length
SAC	Standard	1	0.33±0.07	100±0	0.45±0.12	0.0±0.0	86.7±19.5
		5	0.91±0.12	100±0	0.58±0.12	0.0±0.0	76.8±16.2
		10	1.65±0.31	100±0	0.46±0.12	0.0±0.0	78.7±16.5
		20	4.52±1.43	98±2	0.35±0.11	0.0±0.0	63.2±18.4
		30	9.23±0.96	47±38	0.19±0.04	0.0±0.0	525.0±63.1
		50	10.00±0.00	7±5	0.06±0.02	0.0±0.0	917.0±35.5
	RAD	1	0.24±0.03	100±0	0.62±0.15	0.0±0.0	50.9±6.8
		5	0.42±0.06	100±0	0.42±0.13	0.0±0.0	76.9±12.0
		10	0.82±0.18	100±0	0.30±0.11	0.0±0.0	144.1±25.4
	CURL	1	0.46±0.10	100±0	0.76±0.09	0.0±0.0	77.2±13.1
		5	1.56±0.31	100±0	0.44±0.10	0.2±0.4	84.9±17.7
		10	5.24±1.92	100±0	0.37±0.11	0.0±0.0	88.0±20.3
	SPR	1	2.09±0.81	100±0	0.65±0.13	0.6±1.2	206.4±23.8
		5	3.68±1.68	69±13	0.21±0.09	0.0±0.0	523.9±55.6
		10	10.00±0.00	9±3	0.07±0.04	0.0±0.0	912.1±36.0
	DBC	1	0.44±0.04	100±0	0.44±0.13	0.0±0.0	65.1±13.8
		5	0.99±0.25	100±0	0.34±0.13	0.0±0.0	111.2±22.7
		10	10.00±0.00	2±1	0.13±0.04	0.0±0.0	588.5±63.3
	AE	1	0.42±0.09	100±0	0.78±0.11	0.0±0.0	85.4±22.1
		5	1.04±0.24	100±0	0.64±0.16	0.0±0.0	102.6±22.6
		10	2.03±0.38	100±0	0.55±0.12	1.3±2.6	78.0±17.1
	VAE	1	1.13±0.14	100±0	0.64±0.15	0.0±0.0	75.1±15.6
		5	5.30±0.68	100±0	0.30±0.08	0.0±0.0	81.2±18.2
		10	10.00±0.00	25±17	0.12±0.03	0.4±0.5	834.3±47.5
	SB	1	0.98±0.88	100±0	0.89±0.08	0.0±0.0	130.1±25.0
		5	2.08±0.30	91±17	0.65±0.12	0.0±0.0	117.8±20.5
		10	10.00±0.00	3±2	0.06±0.02	0.2±0.4	980.3±18.0

Table 13. Performance metrics for DreamerV3 agents across image pool sizes. ‘Steps’ indicates the number of environment steps (in millions) required to reach 80% success rate during training. ‘ID Success’ shows the final success rate at the end of 10M training steps when evaluating on the same pool used during training. ‘Length’ indicates the average number of steps required to solve an instance of the puzzle at the last 100 episodes.

Algorithm	Variation	Pool	Steps (M)	ID Success (%)	Length
DreamerV3	Standard	1	0.42±0.06	100±0	83.5±11.8
		5	1.23±0.20	100±0	31.0±3.9
		10	1.44±0.58	100±0	32.9±4.0
		20	3.96±0.61	100±0	27.8±3.5
		30	5.84±0.71	99±1	38.0±6.2
		50	6.62±2.67	87±14	177.8±24.8
	w/o decoder	100	8.18±2.33	29±14	676.1±32.8
		1	1.13±0.12	100±0	36.0±2.5
		5	1.79±0.61	100±0	42.1±3.2
			10	2.57±0.91	100±0

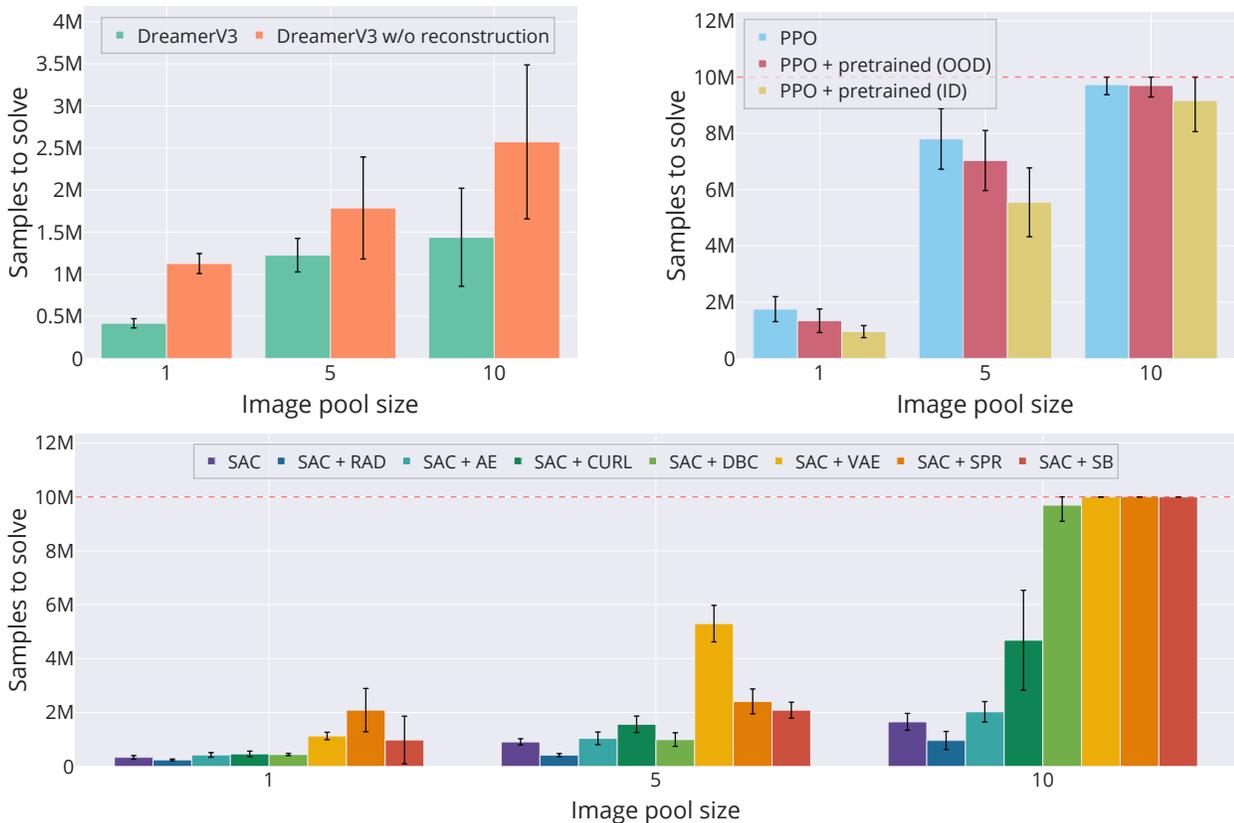


Figure 9. Sample efficiency of different methods across pool sizes (lower is better). SPGym differentiates agents based on their representation learning capabilities. Top left: DreamerV3 variants demonstrate the value of reconstruction learning. Top right: PPO results show benefits of pretraining. Bottom: Comprehensive comparison of SAC variants reveals trade-offs between different representation learning approaches.

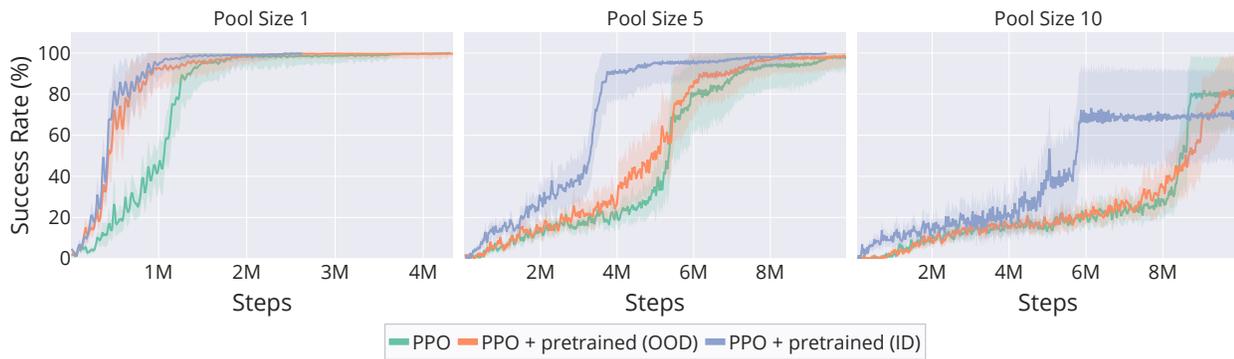


Figure 10. **Learning curves for PPO variants.** Success rate during training for baseline PPO and versions with pretrained encoders across different pool sizes. Shaded regions represent 95% confidence intervals across 5 independent runs.

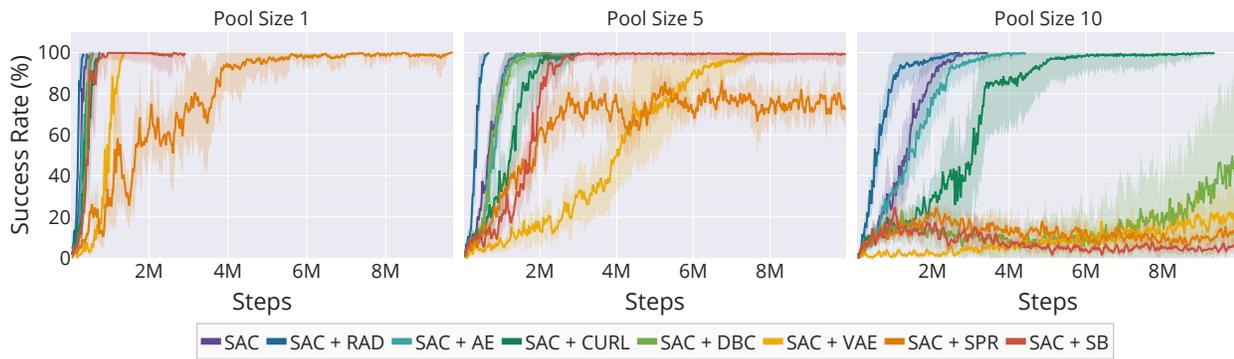


Figure 11. **Learning curves for SAC variants.** Success rate during training for baseline SAC and versions with different representation learning components across different pool sizes. Shaded regions represent 95% confidence intervals across 5 independent runs.

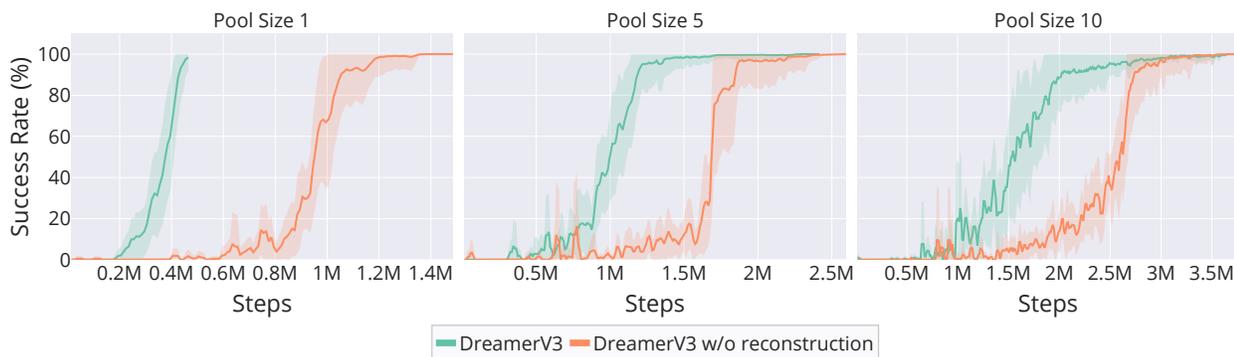


Figure 12. **Learning curves for DreamerV3 variants.** Success rate during training for DreamerV3 with and without decoder across different pool sizes. Shaded regions represent 95% confidence intervals across 5 independent runs.