Disentangling Recognition and Decision Regrets in Image-Based Reinforcement Learning

Alihan Hüyük, A. Ryo Koblitz, Atefeh Mohajeri, Matthew Andrews

Keywords: image-based reinforcement learning, observational overfitting, over-specific representations, under-specific representations, recognition regret, decision regret

Summary

In image-based reinforcement learning (RL), policies usually operate in two steps: first extracting lower-dimensional features from raw images (the "recognition" step), and then taking actions based on the extracted features (the "decision" step). Extracting features that are spuriously correlated with performance or irrelevant for decision-making can lead to poor generalization performance, known as *observational overfitting* in image-based RL. In such cases, it can be hard to quantify how much of the error can be attributed to poor feature extraction vs. poor decision-making. To disentangle the two sources of error, we introduce the notions of *recognition regret* and *decision regret*. Using these notions, we characterize and disambiguate the two distinct causes behind observational overfitting: *over-specific representations*, which include features that are not needed for optimal decision-making (leading to high decision regret), vs. *under-specific representations*, which only include a limited set of features that were spuriously correlated with performance during training (leading to high recognition regret). Finally, we provide illustrative examples of observational overfitting due to both over-specific and under-specific representations in maze environments and the Atari game Pong.

Contribution(s)

1. We define recognition and decision regrets, which disentangle the regret induced by poor recognition policies vs. poor decision policies.
   Context: In image-based RL, most agents first extract features from images and then take actions based on the extracted features. When an RL agent does not perform well (measured by its regret), it is hard to tell whether this is due to a failure to extract features or a failure to plan good actions. Our definitions break down overall regret into two components that attribute it to one of these two failure modes.
2. By analyzing generalization performance through recognition and decision regrets, we characterize over-specific and under-specific representations as two distinct modes of observational overfitting in image-based RL.
   Context: Observational overfitting is a phenomenon in image-based RL that an agent learns to rely on information in the image that does not actually constitute a part of the environment state (like decorative elements in a video game) (Song et al., 2020). Using the notions of recognition and decision regrets, we identify this phenomenon can occur in two distinct forms: learning over-specific representations that include irrelevant features (not informative of the actual state) vs. under-specific representations that only include spurious features (correlated with the actual state but only during training, the correlations do not generalize).
3. We provide illustrative examples of observational overfitting due to both over-specific and under-specific representations in maze environments as well as the Atari game Pong.
   Context: These two environments contrast each other well in terms of their dimensionality, and they are commonly considered when exploring generalization in RL (e.g. Sonar et al., 2021; Taiga et al., 2023).

Disentangling Recognition and Decision Regrets in Image-Based Reinforcement Learning

Alihan Hüyük[1,†], **A. Ryo Koblitz**[2], **Atefeh Mohajeri**[2], **Matthew Andrews**[2]
ahuyuk@seas.harvard.edu,
{ryo.koblitz,atefeh.mohajeri,matthew.andrews}@nokia-bell-labs.com

[1]**Harvard University**
[2]**Nokia Bell Labs**

[†]Work done as an intern at Nokia Bell Labs

Abstract

In image-based reinforcement learning (RL), policies usually operate in two steps: first extracting lower-dimensional features from raw images (the "recognition" step), and then taking actions based on the extracted features (the "decision" step). Extracting features that are spuriously correlated with performance or irrelevant for decision-making can lead to poor generalization performance, known as *observational overfitting* in image-based RL. In such cases, it can be hard to quantify how much of the error can be attributed to poor feature extraction vs. poor decision-making. To disentangle the two sources of error, we introduce the notions of *recognition regret* and *decision regret*. Using these notions, we characterize and disambiguate the two distinct causes behind observational overfitting: *over-specific representations*, which include features that are not needed for optimal decision-making (leading to high decision regret), vs. *under-specific representations*, which only include a limited set of features that were spuriously correlated with performance during training (leading to high recognition regret). Finally, we provide illustrative examples of observational overfitting due to both over-specific and under-specific representations in maze environments and the Atari game Pong.

1 Introduction

In image-based reinforcement learning (RL), it is common for agents to act in two steps by following two sub-policies one after another (e.g. the popular RL package Stable Baselines3, Raffin et al., 2021): (i) First, a *recognition policy* generates representations of input images by extracting lower-dimensional features from high-dimensional pixels. (ii) Then, a *decision policy* generates actions based on those representations (see Figure 1). Naturally, if an agent acting in this way is performing poorly, this would be partly due to poor feature extraction and partly due to poor decision-making, in other words, partly because the recognition policy is suboptimal and partly because the decision policy is suboptimal. Here, identifying which policy is the main source of error in an underperforming RL system can provide crucial information for debugging and improving that system. However, quantifying the individual contribution of each policy to the overall error can be challenging, especially when agents are trained in an end-to-end fashion and both policies are optimized jointly.

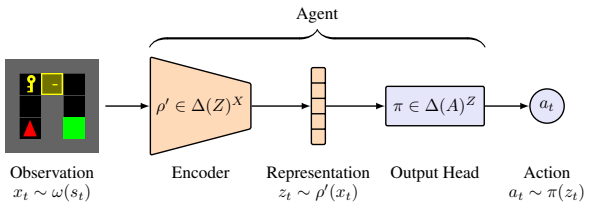

Figure 1: Block diagram of an agent consisting of a *recognition policy* (encoder) followed by a *decision policy* (output head).

Our goal in this paper is to disentangle these two potential sources of error, namely poor recognition vs. poor decision-making. As a motivating example of how such a distinction can play an important role in building RL systems, consider *observational overfitting* typically encountered in image-based RL (Song et al., 2020): Image-based policies often do not generalize well to unseen scenarios because they fail to extract *critical features* that need to be identified in images to be able to make good decisions. Instead, they end up extracting either *irrelevant features* that should not matter when making decisions or *spurious features* that were mistakenly correlated with performance during training. As an example of irrelevant features, a self-driving car might encode in its representations the makes and models of other cars in the traffic although good driving does not require recognizing car makes and models, and consequently, might struggle when encountering a new car with a never-before-seen make and model. As an example of spurious features, a self-navigating robot might memorize landmarks in the background as cues for when to perform certain maneuvers rather than identifying actual obstacles, and consequently, might fail to navigate a new course when the same landmarks are no longer there. In order to avoid observational overfitting, we should be able to tell when a recognition policy fails to extract critical features over irrelevant or spurious features.

We achieve our goal of disentangling the error induced by an agent's recognition policy vs. their decision policy by introducing the notions of *recognition regret* and *decision regret*. Conventionally, *regret* measures the difference between an agent's performance and the best possible performance of any agent. Our regret notions break this conventional definition into two parts: (i) First, relying on representations generated by a given recognition policy as input to decision policies, rather than the original images directly, might already reduce the best performance that can possibly be achieved by any decision policy—the difference is the recognition regret. (ii) Then, a given decision policy might still fall short of achieving the best possible performance of any decision policy—the difference is the decision regret. In other words, the recognition regret measures how much the representations of a given recognition policy bottlenecks the down-stream performance of decision policies while the decision regret measures how much potential performance is missed out by a given decision policy.

Making use of our definitions, we characterize and disambiguate two distinct causes behind observational overfitting: *over-specificity* or *under-specificity* of representations generated by recognition policies. Over-specific representations include "too many" features, often irrelevant ones in addition to critical ones. Since these representations still include the critical features, they do not necessarily limit the downstream performance of decision policies and hence do not incur much recognition regret. However, the presence of irrelevant features might cause decision policies to treat images with similar critical features as if they are unrelated when they do not share the same irrelevant features, which would lead to high decision regret. In contrast, under-specific representations include "too few" features, usually only capturing spurious indicators of performance during training. This might cause decision policies to treat unrelated images as if they are similar to each other just because they feature the same spurious indicators, which would lead to high recognition regret over decision regret.

Contributions Our contributions are three-fold: First, in **Section 3**, we define recognition and decision regrets, which disentangle the regret induced by poor recognition policies vs. poor decision policies. Next, in **Section 4**, we apply these definitions to the problem setting of generalization in RL. Through recognition and decision regrets in terms of generalization performance, we characterize over-specific and under-specific representations as two distinct modes of observational overfitting in image-based RL. Finally, in **Section 5**, we provide illustrative examples of observational overfitting due to both over-specific and under-specific representations in maze environments as well as the Atari game Pong. These examples demonstrate that the notions of recognition and decision regrets can be powerful tools in articulating the cause of observational overfitting in image-based RL applications.

2 Related Work

Representations in Image-Based RL In RL, algorithms for training policies in the form of neural networks generally fall into one of two categories: action-value fitting methods such as Mnih et al. (2015) and policy gradient methods such as Mnih et al. (2016); Schulman et al. (2017). When these

neural networks take images as input, one of the most common architectures used in practice is an encoder consisting of convolutional layers followed by either an action-value head if using an action-value fitting method (e.g. Mnih et al., 2015; Raffin et al., 2021), or followed by a policy head for generating actions and a value head for estimating values if using a policy gradient method (e.g. Cobbe et al., 2019; Raffin et al., 2021). With this type of architecture (an encoder followed by output heads), the intermediary output of the encoder can be viewed as a lower-dimensional representation of the input images. While a lot of previous work has focused on learning "better" representations for various purposes (e.g. Sonar et al., 2021; Raileanu & Fergus, 2021; Dabney et al., 2021; Agarwal et al., 2021; Stooke et al., 2021; Zhang et al., 2022; Eysenbach et al., 2022), our focus in this paper is to *quantify* the quality of an encoder's representations as far as its contribution to the end performance is concerned. We capture the performance loss due to a poor encoder (i.e. a poor recognition policy) through recognition regret. Meanwhile, any additional losses that might be incurred by the subsequent action-value/policy/value heads (i.e. the decision policies) are captured through decision regret.

Others have also asked what constitutes a "good" representation in RL: According to Bellemare et al. (2019); Le Lan et al. (2022), representations should be sufficient in approximating the value of all stationary policies. Meanwhile, Ghosh & Bellemare (2020); Wang et al. (2024) survey different representation learning schemes and evaluate a variety of properties such as stability/robustness, capacity, and redundancy. Different from all this work, we consider the representations extracted by a recognition policy not in isolation but in relation to a specific decision policies, which may have been trained jointly with the recognition policy or subsequently (i.e. recognition regret w.r.t. decision regret). According to our work, representations should not only avoid hindering performance (cf. recognition regret) but also facilitate the training of high-performing decision policies (cf. decision regret). Under-specific representations fail at the former while over-specific representations fail at the latter.

Generalization in RL & Observational Overfitting Training policies that generalize beyond their particular training conditions remain a significant challenge in RL (Zhang et al., 2018a; Packer et al., 2018; Nichol et al., 2018). This is partly because policies may overfit due to a multitude of confounded factors that can be hard to disambiguate from each other. For instance, some of these factors may include the characteristics of the training environment like whether it is deterministic or not (Bellemare et al., 2013; Pinto et al., 2017; Rajeswaran et al., 2017; Machado et al., 2018), the method of exploration during training like how many initial states are sampled (Zhang et al., 2018a;b), or even the hyper-parameters of the training algorithms like the discount factor γ (Jiang et al., 2015).

We focus our attention on observational overfitting specifically as characterized by (Song et al., 2020) because of its close relationship to learned representations in image-based RL. This particular type of overfitting occurs in cases where different environments are all semantically the same with the same state spaces and the same transition dynamics but generate different images/observations. For instance, a game might have the same underlying game mechanics in all of its levels but each level might have its own unique art style (e.g. the popular generalization benchmarks Nichol et al., 2018; Cobbe et al., 2019). In such cases, if a policy fails to learn representations that accurately identify the shared states between different environments, and instead extracts features that are merely cosmetic and differ from one environment to another, then that policy might fail to generalize to new environments that have unseen observation dynamics (Sonar et al., 2021). In this paper, we go one step beyond Song et al. (2020) and describe two distinct modes of observational overfitting: due to over-specificity, where the policy successfully extracts the true state but also features that are irrelevant to transition dynamics, and due to under-specificity, where the policy fails to extract the true state in the first place and ends up relying on spurious features that do not generalize.

3 Defining Recognition and Decision Regrets

So far, when we have stated that good representations should facilitate the training of high-performing policies, we have left the meaning of "high-performing" intentionally vague so that the idea of decomposing regret into recognition and decision components remains broadly applicable. Depending on the application, this might simply mean achieving high value on a given task within operational

constraints. For instance, if we are in a setting limited by compute, learning representations that allow for simpler decision policies would be desirable. Another examples is generalization: If the primary concern is for the decision policies to remain robust to changes in environment dynamics, then we would want our representations to include features that remain critical considerations for decision-making across a wide range of environment conditions.

In this section, we give formal definitions of decision regret and recognition regret for the standard RL setting as the simplest initiation of our idea, where we are only interested in maximizing value for a single task. We also aim to provide an intuitive understanding of these definitions through a worked example in Section 3.1 and empirical examples in Section 3.2. Later in Section 4, we will apply these two concepts to the problem of generalization, which involve training and test environments with different dynamics. This application not only demonstrates how the same decomposition can be applied to a different notion of performance, but also allows us to describe two potential causes behind observational overfitting: over-specific vs. under-specific representations. Finally in Section 5, we will provide illustrative examples of both these causes in action.

Environments Consider partially-observable environments: $\varepsilon = (\sigma, \tau, \omega, r)$ with *state* space S, *action* space A, and *observation* space X, where $\sigma \in \Delta(S)$ is the initial state distribution, $\tau \in \Delta(S)^{S\times A}$ is the *transition dynamics*, $\omega \in \Delta(X)^S$ is the *observation dynamics*, and $r \in \Delta(\mathbb{R})^S$ is the *reward dynamics*. Beginning from an initial state $s_1 \sim \sigma$, at each time step $t \in \{1, 2, \ldots\}$, the environment emits an observation $x_t \sim \omega(s_t)$ based on the current state, gives out a reward $r_t \sim r(s_t)$, and finally transitions into a new state $s_{t+1} \sim \tau(s_t, a_t)$ based on the actions of an agent.

Recognition & Decision Policies For some *representation* space Z, we denote with $\rho \in \Delta(Z)^{Z\times X}$ *recognition policies* and with $\pi \in \Delta(A)^Z$ *decision policies*. Beginning from an initial representation $z_0 \in Z$, an agent following these policies first generates an updated representation $z_t \sim \rho(z_{t-1}, x_t)$ based on the most recent observation according to the recognition policy and then takes an action $a_t \sim \pi(z_t)$ based on the updated representation according to the decision policy. For agents with no memory, the recognition policy would be a function of the most recent observation only (i.e. an encoder) such that $z_t \sim \rho(z_{t-1}, x_t) \doteq \rho'(x_t)$ where $\rho' \in \Delta(Z)^X$. Given a pair of recognition and decision policies (ρ_0, π_0) and a discount factor $\gamma \in (0, 1)$, their (discounted) *value* in environment ε is

$$V_\varepsilon(\rho_0, \pi_0) = \mathbb{E}_{\varepsilon,\rho_0,\pi_0}\left[\textstyle\sum_t \gamma^t r_t\right] \tag{1}$$

Recognition & Decision Regrets Having defined value, we can now quantify the suboptimality of a given policy (ρ_0, π_0) in environment ε through its *regret*, which is simply the gap between its value and the *optimal value* V^*_ε (i.e. how much of the attainable value is missed out by the given policy):

$$R_\varepsilon(\rho_0, \pi_0) \doteq V^*_\varepsilon - V_\varepsilon(\rho_0, \pi_0) \doteq \max_{\rho,\pi} V_\varepsilon(\rho, \pi) - V_\varepsilon(\rho_0, \pi_0) \tag{2}$$

We decompose this regret into two parts, *recognition regret* and *decision regret* respectively:

$$R^{\text{rec}}_\varepsilon(\rho_0) \;\doteq\; \max_{\rho,\pi} V_\varepsilon(\rho, \pi) - \max_\pi V_\varepsilon(\rho_0, \pi) \tag{3}$$

$$R^{\text{dec}}_\varepsilon(\rho_0, \pi_0) \;\doteq\; \max_\pi V_\varepsilon(\rho_0, \pi) - V_\varepsilon(\rho_0, \pi_0) \tag{4}$$

such that $R = R^{\text{rec}} + R^{\text{dec}}$. Intuitively, the recognition regret measures how much ρ_0 alone limits the value attainable by subsequent decision policies. While generally, the value can be as high as the optimal value $V^* = \max_{\rho,\pi} V(\rho, \pi)$, when conditioned on ρ_0 specifically, it becomes bounded by $\max_\pi V(\rho_0, \pi) \leq V^*$ (and the difference is defined as R^{rec}). Meanwhile, the decision regret measures how much of the value still attainable after conditioning on ρ_0 is missed out by π_0. Although the value can still be as high as $\max_\pi V(\rho_0, \pi)$ under ρ_0, pairing ρ_0 with π_0 leads only to a value of $V(\rho_0, \pi_0)$ (and the difference is defined as R^{dec}).

3.1 A Worked Example

To illustrate our definitions, we consider the simple 3-state environment with $S = \{0, 1, 2\}$ shown in Figure 2. Only state 2 produces a reward so the optimal policy takes action 1 in states 0 and 1,

and action 0 in state 2, since this ensures we achieve reward 1 in half of the time steps. If the initial state is 0, the optimal value in this environment is given by v_0 in the equations $v_0 = \gamma v_1$, $v_1 = \gamma v_2$, $v_2 = 1 + \gamma v_1$, that is $V^* = \gamma^2 + \gamma^4 + \cdots$, which equals 4.263 when the discount factor is $\gamma = 0.9$.

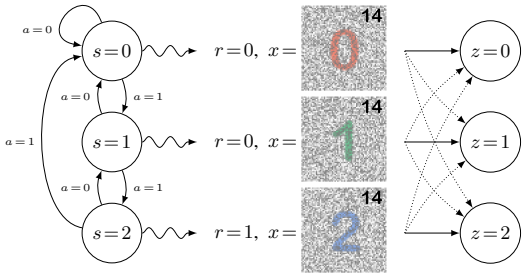

Figure 2: States **(left)**, observations **(center)**, and representations **(right)** for our 3-state example at time step $t = 14$.

The observation space X is a union of image spaces $\cup_{i \in S, t \in \{1,2,\ldots\}} I_{it}$, where I_{it} is a set of noisy representations of the number i, with a clear black representation of the time t in the top right corner. The colors of the integers in I_{0t}, I_{1t}, I_{2t} are red, green, blue respectively. When $s_t = i$, observation x_t is an image in I_{it} selected randomly.

The Recognition Policy The recognition policy ρ is a function of the most recent observation only and is defined by a mapping $\rho'(x)$ from each image x to the representation space $Z = \{0, 1, 2\}$ (hence $Z = S$). The optimal recognition policy "should" recognize the depiction of s_t in the image (and ignore the time in the top right corner). However, the noise in the images means that I_{0t}, I_{1t}, I_{2t} are hard to distinguish and so consider a suboptimal policy ρ'_0 defined for some error parameter δ. When $s_t = i$, with probability $1 - \delta$, $\rho'_0(x_t) = i$ (i.e. the recognition policy is correct), and with probability δ, $\rho'_0(x_t)$ is uniformly distributed across Z (i.e. the recognition policy makes an arbitrary choice).

When δ is small, the optimal decision policy is still to take action 1 in states 0 and 1, and action 0 in state 2 (since we always want to get to state 2 as soon as possible). However, with probability δ, $\rho'_0(x_t)$ is a random element from Z due to incorrect recognition, and when that happens the action is 1 with probability $2/3$ and 0 with probability $1/3$. The value in the environment is now given by v_0 in the equations $v_0 = \gamma v_0(\delta/3) + \gamma v_1(1 - \delta/3)$, $v_1 = \gamma v_0(\delta/3) + \gamma v_2(1 - \delta/3)$, $v_2 = 1 + \gamma v_0(2\delta/3) + \gamma v_1(1 - 2\delta/3)$. When $\gamma = 0.9$ and $\delta = 0.1$, the solution to these equations is $(v_0, v_1, v_2) = (3.992, 4.451, 4.978)$ hence the *recognition regret* $R^{\text{rec}}(\rho_0) = 4.263 - 3.992 = 0.271$.

The Decision Policy Suppose we now couple this suboptimal recognition policy with a δ-optimal decision policy π_0, meaning one that based on the state recognition takes the optimal action with probability $1 - \delta$ and a random action with probability δ. If we combine the recognition and decision errors, if the true state is 0 then π_0 takes the wrong action 0 with probability $\delta/3 + \delta/2$, if the true state is 1 then π_0 takes the wrong action 0 with probability $\delta/3 + \delta/2$, and if the true state is 2 then π_0 takes the wrong action 1 with probability $2\delta/3 + \delta/2$. The value in the environment is now given by v_0 in the equations $v_0 = \gamma v_0(\delta/3 + \delta/2) + \gamma v_1(1 - \delta/3 - \delta/2)$, $v_1 = \gamma v_0(\delta/3 + \delta/2) + \gamma v_2(1 - \delta/3 + \delta/2)$, $v_2 = 1 + \gamma v_0(2\delta/3 + \delta/2) + \gamma v_1(1 - 2\delta/3 - \delta/2)$. When $\gamma = 0.9$ and $\delta = 0.1$, the solution to these equations is $(v_0, v_1, v_2) = (3.680, 4.126, 4.666)$ hence the *decision regret* $R^{\text{dec}}(\rho_0, \pi_0) = 3.992 - 3.680 = 0.312$.

3.2 Empirical Examples

To illustrate our definitions further, we also provide empirical examples where we intentionally construct suboptimal decision policies and suboptimal recognition policies and highlight how the values of R^{rec} and R^{dec} differ between the two scenarios we construct.

Setup We consider two image-based environments: *Minigrid*, which is a maze environment where the agent needs to navigate around walls to reach some target (Chevalier-Boisvert et al., 2023), and *Pong*, which is an Atari game where the agent aims to score points against a computer opponent (Towers et al., 2023). In both environments, we consider policies with the same type of architecture, where the recognition policy, $z = \rho_{\text{CNN}}(x)$, is given by a convolutional neural network (CNN) and the decision policy, $a \sim \pi_{\text{MLP}}(z)$, is given by a multi-layer perceptron (MLP), and we train these policies using the PPO algorithm (Schulman et al., 2017) from the skrl package (Serrano-Muñoz et al., 2023). When presenting results, we shift and scale rewards so that the regret is at most one. Details regarding the setup can be found in the supplementary material.

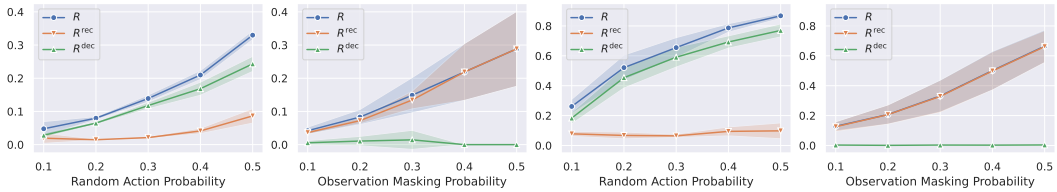

(a) Minigrid **(b) Pong**

Figure 3: *Recognition and Decision Regrets under Randomized Actions and Masked Observations.* Both scenarios lead to regret. For random actions, it is mostly decision regret, indicating: Even when forced to occasionally take random actions, agents still learn useful representations of their environment (to be able to optimize their non-random actions). For masked observations, it is recognition regret, indicating: Partially masked observations fundamentally limit the agent's ability to optimize their actions regardless of their decision policy.

Suboptimal Decision Policies We construct agents with *suboptimal* decision policies by restricting them to occasionally take random actions. First, we modify the network π_{MLP}, denoting with $\tilde{\pi}$ the resulting decision policy, such that

$$a \sim \tilde{\pi}(z) = \begin{cases} \text{Uniform}(A) & \text{with probability} \quad p \\ \pi_{\text{MLP}}(z) & \text{with probability} \quad 1-p \end{cases} \tag{5}$$

where $p \in [0, 1]$ is the probability of taking a forced random action and $\text{Uniform}(A)$ is the uniform distribution over A. Then, we train the weights of the two networks, ρ_{CNN} and π_{MLP}, as usual via PPO.

In this scenario, the occasional random actions would inevitably induce some regret $R(\rho_{\text{CNN}}, \tilde{\pi})$. However, maximizing the rewards when actions are not forced to be taken randomly would still require finding a near-optimal ρ_{CNN} and π_{MLP} since non-random actions are still generated by $\rho_{\text{CNN}} \circ \pi_{\text{MLP}}$. Therefore, we would expect that $V(\rho_{\text{CNN}}, \pi_{\text{MLP}}) \approx V^*$ and $\pi_{\text{MLP}} \approx \operatorname{argmax}_\pi V(\rho_{\text{CNN}}, \pi)$, and hence, $R^{\text{rec}}(\rho_{\text{CNN}}) = V^* - \max_\pi V(\rho_{\text{CNN}}, \pi) \approx 0$. In other words, the occasional random actions should not induce much recognition regret but rather induce decision regret.

Figure 3 confirms this intuition, where we plot R, R^{rec}, R^{dec} of $(\rho_{\text{CNN}}, \tilde{\pi})$ for varying p on the left-side panels. Note that we scale rewards so that $V^* = 1$, hence regret can be at most be one. Indeed, an increasing probability of random actions leads to increasing regret as well which is predominantly decision regret in both *Minigrid* and *Pong*. In these panels, plotting R^{rec} and R^{dec} requires estimating $\max_\pi V(\rho_{\text{CNN}}, \pi)$, which we achieve by fixing the weights of ρ_{CNN}, setting $p = 0$ to remove any restrictions on π_{MLP} so that $\tilde{\pi} = \pi_{\text{MLP}}$, and re-training the weights of π_{MLP}.

Suboptimal Recognition Policies We construct agents with *suboptimal* recognition policies by randomly masking parts of the images/observations before feeding them forward to an optimal agent. First, we train the weights of ρ_{CNN} and π_{MLP} without any modifications. Then, letting $X = \mathbb{R}^{W \times H \times C}$ be the space of W-by-H images with C channels, we modify the (pre-trained) network ρ_{CNN}, denoting with $\tilde{\rho}$ the resulting recognition policy, such that

$$z = \tilde{\rho}(x) = \rho_{\text{CNN}}(x \odot m), \quad m_{ijk} = \begin{cases} 0 & \text{w/ prob.} \quad p \\ 1 & \text{w/ prob.} \quad 1-p \end{cases} \tag{6}$$

where $p \in (0, 1)$ is the probablity of masking an individual pixel-channel, $m \in \{0, 1\}^{W \times H \times C}$ is a randomly sampled masking vector, and $\odot$ denotes element-wise multiplication.

In this scenario, the partially masked observations would induce some regret $R(\tilde{\rho}, \pi_{\text{CNN}})$ as before. However this time, we would expect most of that regret to be recognition regret since switching from ρ_{CNN} to $\tilde{\rho}$ would limit the performance of all decision policies π not just π_{CNN} specifically.

Again, Figure 3 confirms this intuition, where we plot R, R^{rec}, R^{dec} of $(\tilde{\rho}, \pi_{\text{MLP}})$ for varying p on the right-side panels: An increasing masking probability leads to increasing regret that is predominantly recognition regret in both *Minigrid* and Pong. Just like before, plotting R^{rec} and R^{dec} in these panels requires estimating $\max_\pi V(\tilde{\rho}, \pi)$, which we achieve by fixing the weights of ρ_{CNN} and the value of p so that the recognition policy $\tilde{\rho}$ is completely fixed, and then retraining the weights of π_{MLP}.

4 Practical Application: Generalization in terms of Recognition and Decision Regrets

In practice, it is not always possible to directly measure the decision, recognition, or overall regret of policies as these metrics require knowing optimal policies under various constraints. Often, our goal is to find such optimal policies in the first place. For instance, estimating $\max_\pi V(\rho_0, \pi)$ would require us to optimize a decision policy given a recognition policy. Suppose we happen to have a method to do so with the result being $\hat{\pi}^* \approx \operatorname{argmax}_\pi V(\rho_0, \pi)$. Then, we can estimate the decision regret as $R^{\text{dec}} \approx V(\rho_0, \hat{\pi}^*) - V(\rho_0, \pi_0)$, which is indeed what we did in our empirical examples in Section 3.2. This is a sound approach if π_0 is obtained through a different process, perhaps it is optimized under compute constraints or it belongs to an interpretable class of policies, then comparing π_0 to an unconstrained or uninterpretable baseline $\hat{\pi}^*$ would be informative. However, in applications where performance is the primary concern, it is likely the case that the decision policy we are interested in analyzing, i.e. π_0, is already our best estimate of the optimal decision policy, i.e. $\hat{\pi}^*$. This would render our estimate of the decision regret meaningless as it would be equal to zero identically.

However, when measuring *generalization performance*, that is the performance of policies optimized in some training environment but evaluated in some other test environment, then it can be possible to estimate recognition and decision regrets in a meaningful way (essentially by computing regret via policies optimized in the test environment to analyze policies that are optimized in the training environment). In this section, we develop this idea as a practical application of our definitions in Section 3. First, we introduce the notion of *generalization regret* to quantify how much performance is missed out in the test environment due to suboptimal generalization from the training environment. Then, similar to our approach in the previous section, we decompose generalization regret into recognition and decision components. These two components allow us to characterize over-specific vs. under-specific representations, which we exemplify later in Section 5.

Generalization Regret Instead of a single environment ε, we now consider two different environments: a *training environment* $\varepsilon^{\text{train}}$ and a *test environment* $\varepsilon^{\text{test}}$, which we denote together as $\mathcal{E} = (\varepsilon^{\text{train}}, \varepsilon^{\text{test}})$. For instance, if we are interested in observational overfitting, these two environments might differ in terms of their observation dynamics ω (but not τ). Moreover, suppose that we are given a learning algorithm $\mathcal{L}$ for training policies such that $\mathcal{L}(\varepsilon) \doteq (\mathcal{L}_\rho(\varepsilon), \mathcal{L}_\pi(\varepsilon)) \approx \operatorname{argmax}_{\rho,\pi} V_\varepsilon(\rho, \pi)$ and $\mathcal{L}_\pi(\varepsilon|\rho_0) \approx \operatorname{argmax}_\pi V_\varepsilon(\rho_0, \pi)$. Then, the generalization performance of policies trained in $\varepsilon^{\text{train}}$ via $\mathcal{L}$ but tested in $\varepsilon^{\text{test}}$ is given by $V_{\varepsilon^{\text{test}}}(\mathcal{L}(\varepsilon^{\text{train}}))$. Similar to regret in Section 3, we can consider this performance relative to the best possible test performance, which we call the *generalization regret*:

$$GR_{\mathcal{E}}(\mathcal{L}) = V^*_{\varepsilon^{\text{test}}} - V_{\varepsilon^{\text{test}}}(\mathcal{L}(\varepsilon^{\text{train}})) \tag{7}$$

Estimating this metric still requires knowing the optimal test policy $\operatorname{argmax}_{\rho,\pi} V_{\varepsilon^{\text{test}}}(\rho, \pi)$. However, unlike the case with regular regret, we can now obtain a meaningful estimate of GR by using policies $\mathcal{L}(\varepsilon^{\text{test}})$ trained directly in the test environment as a proxy for the optimal test policy (since these policies are now distinctly different from the policies $\mathcal{L}(\varepsilon^{\text{train}})$ trained originally). This results in the following estimation of the generalization regret:

$$\widehat{GR}_{\mathcal{E}}(\mathcal{L}) = V_{\varepsilon^{\text{test}}}(\mathcal{L}(\varepsilon^{\text{test}})) - V_{\varepsilon^{\text{test}}}(\mathcal{L}(\varepsilon^{\text{train}})) \tag{8}$$

which measures how much the test performance could have been improved if it were to be possible to train policies directly in the test environment. Notice the subtle difference between this and GR: The latter measures how much of the test performance is missed out by training policies only in the training environment. When the policies trained in the training environment are already generalizable to the test environment, that is when $GR = 0$, for sensible algorithms $\mathcal{L}$, $\widehat{GR}$ should also be zero, meaning no performance gain should be possible by re-training policies in the test environment.

Relationship with Generalization Error Notably, the generalization regret is closely related to the *generalization error* as previously defined by Zhang et al. (2018a):

$$GE_{\mathcal{E}}(\mathcal{L}) = V_{\varepsilon^{\text{train}}}(\mathcal{L}(\varepsilon^{\text{train}})) - V_{\varepsilon^{\text{test}}}(\mathcal{L}(\varepsilon^{\text{train}})) \tag{9}$$

which measures how much performance is retained when the policies trained in the training environment are rolled out in the test environment. Both definitions are ultimately concerned with the same generalization performance given by $V_{\varepsilon^{\text{test}}}(\mathcal{L}(\varepsilon^{\text{train}}))$. The difference is that, while GE measures this performance relative to the training performance achieved in the training environment, $\widehat{GR}$ measures it relative to the training performance that could have been achieved in the test environment. However, our definition $\widehat{GR}$, being a type of regret, lends itself to the same decomposition as in Section 3 into recognition and decision components, which is not immediately possible with GE.

Recognition & Decision Generalization Regrets Similar to Section 3, consider the intermediary case when only the decision policy has access to the test environment and the test value that would have been achieved then, which can be written as

$$V_{\varepsilon^{\text{test}}}(\ \ \rho = \mathcal{L}_\rho(\varepsilon^{\text{train}})\ ,\ \ \pi = \mathcal{L}_\pi(\varepsilon^{\text{test}}|\mathcal{L}_\rho(\varepsilon^{\text{train}}))\ \ ) \tag{10}$$

Then, we decompose $\widehat{GR}$ with respect to this intermediary value as follows:

$$\widehat{GR}^{\text{rec}}_{\mathcal{E}}(\mathcal{L}) = V_{\varepsilon^{\text{test}}}(\mathcal{L}(\varepsilon^{\text{test}})) - (10) \tag{11}$$

$$\widehat{GR}^{\text{dec}}_{\mathcal{E}}(\mathcal{L}) = (10) - V_{\varepsilon^{\text{test}}}(\mathcal{L}(\varepsilon^{\text{train}})) \tag{12}$$

Intuitively, $\widehat{GR}^{\text{dec}}$ is the hypothetical performance that could be gained from re-training $\mathcal{L}_\pi(\varepsilon^{\text{train}})$ directly in the test environment, and $\widehat{GR}^{\text{rec}}$ is the further performance gain of re-training $\mathcal{L}_\rho(\varepsilon^{\text{train}})$ in addition to $\mathcal{L}_\pi(\varepsilon^{\text{train}})$ (note that $\widehat{GR} = \widehat{GR}^{\text{rec}} + \widehat{GR}^{\text{dec}}$ is the total gain). A worked example, similar to the one in Section 3.1, can be found in the supplementary material.

Overfitting Finally, we are ready to characterize overfitting as well as different modes of overfitting using all the metrics we have defined so far, namely R, $\widehat{GR}$, $\widehat{GR}^{\text{rec}}$, and $\widehat{GR}^{\text{dec}}$: We say that policies $\rho_0, \pi_0 = \mathcal{L}(\varepsilon^{\text{train}})$ have *overfitted* if the generalization regret $GR_{\mathcal{E}}(\mathcal{L}) = R_{\varepsilon^{\text{test}}}(\mathcal{L}(\varepsilon^{\text{train}}))$ is high despite a relatively low training regret $R_{\varepsilon^{\text{train}}}(\mathcal{L}(\varepsilon^{\text{train}}))$. This is said to be a case of *observational overfitting* if the training and test environments share the same dynamics except for their observation dynamics such that $\varepsilon^{\text{train}} = (\sigma, \tau, \omega^{\text{train}}, r)$ and $\varepsilon^{\text{test}} = (\sigma, \tau, \omega^{\text{test}}, r)$. When overfitting occurs,

- We say that representations generated by the recognition policy ρ_0 are *under-specific* if $\widehat{GR}^{\text{rec}}$ is high but $\widehat{GR}^{\text{dec}}$ is relatively low. Intuitively, under-specific representations include "too few" features which do not contain enough information to perform the task at hand (for instance, a blind recognition policy that outputs the same representation for all observations, $\rho'(x_t) = \emptyset$, or an overfitted recognition policy that only learned to extract spurious features). This limits the performance of all decision policies, leading to high $\widehat{GR}^{\text{rec}}$.
- We say that representations generated by the recognition policy ρ_0 are *over-specific* if $\widehat{GR}^{\text{dec}}$ is high but $\widehat{GR}^{\text{rec}}$ is relatively low. Intuitively, over-specific representations include "too many" features which still contain enough information to perform the task at hand but also contain irrelevant information, for instance if the recognition policy has underfitted. This does not necessarily limit the performance that can be achieved by subsequent decision policies, leading to low $\widehat{GR}^{\text{rec}}$, but it may cause a decision policy to take different actions in situations that are semantically the same (i.e. the same state s_t) just because the irrelevant information differs (i.e. different representations z_t), leading to high $\widehat{GR}^{\text{dec}}$. Here, it is the decision policy that has overfitted, due to the over-specific representations generated by an underfitted recognition policy.

In short, under-specific representations may be missing the "signal" in observations and capture only correlated noise, whereas over-specific representations may capture noise (either correlated or uncorrelated) along with the signal.

5 Illustrative Examples

In this section, we give concrete examples of observational overfitting due to both under-specific representations and over-specific representations in two image-based environments: *Minigrid* in Section 5.1, which is a collection of maze environments, and the Atari game *Pong* in Section 5.2.

These environments contrast each other well in terms of their scale.[1] *Minigrid* outputs 7-by-7 images with three channels corresponding to object types, object states, and color identifiers. Meanwhile, *Pong*, after some post-processing, outputs 164-by-82 images that are grayscale and we stack the four most recent images to account for memoryless agents. In addition to these examples, we also highlight model selection as a potential use case of analyzing the specificity of representations in Section 5.3.

As in Section 3.2, we train agents using the PPO algorithm. In all our experiments, the agents consist of a CNN that acts as the recognition policy followed by an MLP that acts as the decision policy (and followed by another parallel MLP that acts as the value head for the PPO algorithm). Details regarding the training setup can be found in the supplementary material.

5.1 Examples in Maze Environments

Setup We consider simple mazes given on a 3-by-3 grid, where the agent needs to navigate to some goal square. We consider different maze configurations, but in each configuration, the goal is separated from the agent by a locked door, and to be able to reach the goal, the agent needs to first pick up a key and unlock the door. With uniform probability, the door can either be at (i) the *North*-end, (ii) the *Center*, or (iii) the *South*-end of the parting wall between the goal and the agent (see Figure 4).

(a) Training Configurations

(b) Test Configurations

Figure 4: *Possible Maze Configurations.* The agent is the red triangle, and the goal is the green square. For training, the key-door color is determined based on the door location, but for testing, the key-door color is always gray (making it both an irrelevant and a spurious feature).

To be able to create agents with over-specified or under-specified representations, our training and testing environments need to have *irrelevant* and *spurious* features. As the irrelevant feature, we introduce color: For the maze configurations in the training environment, the key-door pair is one of three colors: *Red*, *Green*, or *Blue*. In the test environment, the key-door pair is always *Gray*, which is never is never seen during training. These colors are purely visual and they have no impact on how the environment functions, making them an irrelevant feature to the completion of the maze. We turn the key-door color into a spurious feature as well by making it correlated with the door location during training: The north door is always red, the center door is always green, and the south door is always blue in the training configurations. Therefore, knowing the key-door color alone is enough to be able to complete mazes by memorizing the solution to the maze corresponding to each color—but only during training (since the key-door color is always gray during testing).

Agents We construct different agents by applying different *pre-filters* to input images as part of the recognition policy before feeding them forward to the same network architecture. This technique allows us to freely control the properties of an agent's recognition policy and illustrate cases of over-specific vs. under-specific representations. We primarily consider three agents/pre-filters (see the corresponding column in Table 1, results for additional agents can be found in the supplementary material): (i) *Identity* makes no modifications to the input image. This means the key-door color is preserved as an irrelevant feature, making *Identity* susceptible to learning over-specific representations. (ii) *HideDoor* hides the location of the true door by placing false doors in the maze but in the same color as the original door. This means *HideDoor* can only rely on the key-door color as a spurious feature to navigate during training, which may lead to learning *under-specific representations*. (iii) *HideColors* re-colors the key and the door as gray, the same color as the keys and the door in the test environment. This should make *HideColors* easily generalizable to the test environment.

Results Table 1 shows the regret and the generalization regret of each agent. All agents incur almost zero regret during training but only *HideColors* generalizes well to the test environment. *Identity* has a

[1] Besides the scale contrast, these environments are commonly considered when exploring generalization in RL (e.g. Sonar et al., 2021; Taiga et al., 2023). Moreover, they offer organically-generated images whereas the literature on observational overfitting tends to rely mostly on synthetically-generated images, including Song et al. (2020), which introduced the concept.

Table 1: All agents incur almost zero regret during training but only *HideColors* generalizes well to the test environment. *Identity* fails to generalize because its representations are sensitive to irrelevant features (notice the low similarity between the representations of mazes with the same door location but different key-door colors, especially between the unseen gray color and the others). This is characterized by high $\widehat{GR}^{\text{dec}}$ (over-specificity). *HideDoor* fails to generalize because its recognition policy has "overfitted" to extracting spurious indicators of optimal action trajectories (i.e. the key-door color). This is characterized by high $\widehat{GR}^{\text{rec}}$ (under-specificity).

Agent	**Pre-Filter**	**Performance**	**Representation Similarity**		
			All	By Color	By Door Loc.
Identity (***Over-Specific***)	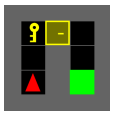	$R = 0.002\ (0.003)$ $\widehat{GR} = 0.911\ (0.171)$ $\widehat{GR}^{\text{rec}} = 0.264\ (0.388)$ $\widehat{GR}^{\text{dec}} = 0.647\ (0.365)$	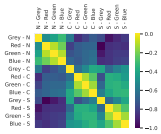	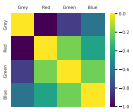	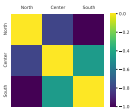
HideColors (***Ideal***)		$R = 0.000\ (0.000)$ $\widehat{GR} = 0.000\ (0.000)$ $\widehat{GR}^{\text{rec}} = 0.000\ (0.000)$ $\widehat{GR}^{\text{dec}} = 0.000\ (0.000)$	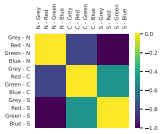	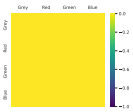	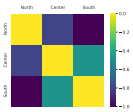
HideDoor (***Under-Specific***)		$R = 0.070\ (0.126)$ $\widehat{GR} = 0.837\ (0.157)$ $\widehat{GR}^{\text{rec}} = 0.607\ (0.227)$ $\widehat{GR}^{\text{dec}} = 0.230\ (0.217)$	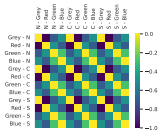	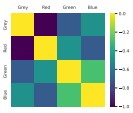	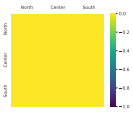

higher $\widehat{GR}^{\text{dec}}$ than $\widehat{GR}^{\text{rec}}$ indicating that its representations are over-specific. This is evidenced by the dis-similarity of its representations across different maze configurations. The similarity matrices in Table 1 plot the average Euclidean distance between the representations encountered while solving one maze vs. another: The closer the representations of two mazes are, the more similar the agent "thinks" those two mazes are to each other. Representations being similar predominantly on the diagonal indicates that *Identity* regard all mazes to be largely independent of each other (even more so with *Gray* mazes as they are unseen during training). In contrast, *HideColors* appropriately regards all mazes with the same door location as exactly the same maze regardless of the key-door color, which is ideal. Meanwhile, *HideDoor* has a higher $\widehat{GR}^{\text{rec}}$ than $\widehat{GR}^{\text{dec}}$ indicating that its representations are under-specific (looking at the similarity matrices, it only differentiates between mazes with different colors).

5.2 Examples in the Atari Game Pong

Setup We consider a deterministic version of Pong and modify the game so that it is reset as soon as the first point is scored. Similar to *Minigrid*, creating agents with over/under-specific representations requires us to introduce irrelevant and spurious features to the game. As the spurious features, we add a frame counter on the top of the game field in the shape of a progress bar (see *Identity* in Table 2). Since the game is completely deterministic, it should be possible to play the game successfully based only on this counter, making it a spurious feature correlated with the optimal action trajectories. However, such a strategy should not generalize to a new initial state (as in the framework of Zhang et al., 2018a). In our experiments, we will vary the initial state by taking a random number of null actions at the start of the game without actually increasing the frame counter.

As the irrelevant features, we generate synthetic images and concatenate them with the original images from the game. We refer to these synthetic images as distractions. They are generated via deterministic functions of the game state and hence contain exactly the same information as the original game screen. Whenever the game is reset, we pick one of two possible generating functions at random, making distractions analogous to decorative elements in other games that change from one level to another (as in the framework of Song et al., 2020). Agents that focus on distractions rather than the proper game field should not generalize to a different set of distractions. In our experiments, we will vary the set of two possible generating functions from training to test time.

Table 2: All agents incur almost zero regret during training. However, *NoCounter* fails to generalize to different distractions because of its over-specific representations that include irrelevant features extracted from distractions. Meanwhile, *JustCounter* fails to generalize to unseen initial states because of its under-specific representations that only include the frame count as a feature, which is a spurious feature since the game is deterministic.

Agent	**Pre-Filter**	**Training Performance**	**Generalization Performance**	
			Unseen Initial States (cf. Zhang et al., 2018a)	Different Distractions (cf. Song et al., 2020)
Identity		$R = 0.001$ (0.002)	$\widehat{GR} = 0.655$ (0.117) $\widehat{GR}^{\text{rec}} = 0.003$ (0.003) $\widehat{GR}^{\text{dec}} = 0.652$ (0.119)	$\widehat{GR} = 0.699$ (0.353) $\widehat{GR}^{\text{rec}} = 0.002$ (0.002) $\widehat{GR}^{\text{dec}} = 0.696$ (0.354)
NoCounter	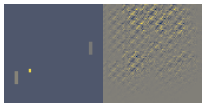	$R = 0.001$ (0.002)	***(Generalizes)*** $\widehat{GR} = 0.001$ (0.002) $\widehat{GR}^{\text{rec}} = 0.000$ (0.000) $\widehat{GR}^{\text{dec}} = 0.001$ (0.002)	***(Over-Specific)*** $\widehat{GR} = 0.642$ (0.263) $\widehat{GR}^{\text{rec}} = 0.001$ (0.002) $\widehat{GR}^{\text{dec}} = 0.641$ (0.262)
JustGameField		$R = 0.006$ (0.008)	***(Generalizes)*** $\widehat{GR} = 0.001$ (0.002) $\widehat{GR}^{\text{rec}} = 0.001$ (0.002) $\widehat{GR}^{\text{dec}} = 0.000$ (0.001)	***(Generalizes)*** $\widehat{GR} = 0.006$ (0.008) $\widehat{GR}^{\text{rec}} = 0.001$ (0.001) $\widehat{GR}^{\text{dec}} = 0.005$ (0.008)
JustCounter		$R = 0.023$ (0.015)	***(Under-Specific)*** $\widehat{GR} = 0.954$ (0.006) $\widehat{GR}^{\text{rec}} = 0.737$ (0.152) $\widehat{GR}^{\text{dec}} = 0.217$ (0.156)	***(Generalizes)*** $\widehat{GR} = 0.023$ (0.015) $\widehat{GR}^{\text{rec}} = 0.002$ (0.003) $\widehat{GR}^{\text{dec}} = 0.021$ (0.016)

Agents Similar to before, we construct different agents via different pre-filters (see the corresponding column in Table 2): (i) *Identity* makes no modifications to the input image. (ii) *NoCounter* hides the frame counter keeping the distractions. This should leave the agent susceptible to learning over-specific representations that include irrelevant features extracted from distractions. (iii) *JustCounter* hides everything but the game counter. This should leave the agent susceptible to learning under-specific representations that only include a spurious feature, the frame count. (iv) *JustGameField* hides both the frame counter (spurious) and any distractions (irrelevant), which is ideal.

Results Table 2 shows the regret and the generalization regret, both for generalization to unseen initial states and different distractions. All agents incur almost zero regret during training. However, *NoCounter* does not generalize to different distractions. This is because the representations learned by *NoCounter* are over-specific (high $\widehat{GR}^{\text{dec}}$ over $\widehat{GR}^{\text{rec}}$). Meanwhile, *JustCounter* does not generalize to unseen initial states. This is because the representations learned by *JustCounter* are under-specific (high $\widehat{GR}^{\text{rec}}$ over $\widehat{GR}^{\text{dec}}$).

5.3 Use Case: Model Selection

Analyzing the recognition and decision generalization regrets of RL models can help make informed decisions about how to improve generalization. If generalization performance is poor due to over-specific representations, we may switch to an algorithm that is capable of learning invariances within the observation space (e.g. Sonar et al., 2021). Conversely, if our representations are under-specific, we may introduce additional loss terms to encourage the model to extract task-critical features.

As a concrete example, we consider one of the simplest adjustments to make to the model

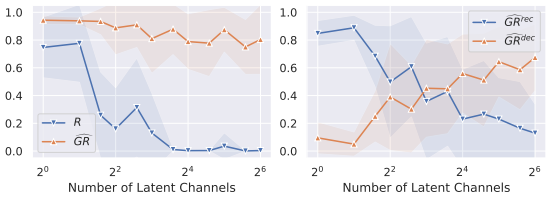

Figure 5: *Regret and Generalization Regrets of Identity in Maze Environments.* Reducing the number of latent channels prevents over-specific representations and lowers $\widehat{GR}^{\text{dec}}$. However, this also results in the recognition policy overfitting to the key-door color, which increases $\widehat{GR}^{\text{rec}}$, ultimately leading to no improvements in terms of the overall generalization regret $\widehat{GR}$.

architecture: changing the representation dimensionality. We consider *Identity* from the examples in maze environments, whose recognition policy outputs a latent representation with $2^6 = 64$ channels. Observing that *Identity* has a high $\widehat{GR}^{\text{dec}}$ relative to $\widehat{GR}^{\text{rec}}$, we conclude that generalization fails due to over-specific representations. To discourage the model from extracting irrelevant features, we decide to reduce the number of latent channels. We can see in Figure 5 that our intervention is successful in lowering $\widehat{GR}^{\text{dec}}$, but the model then ends up overfitting to spurious features like the key-door color instead, resulting in high $\widehat{GR}^{\text{rec}}$ and no overall improvement in terms of the generalization regret. Now that we know our setting suffers both from irrelevant and spurious features, we can start investigating alternative ways to address either one of these problematic features, adopting an appropriate number of latent channels to the issue we target. For instance, we can also see in Figure 5 that the increase in $\widehat{GR}^{\text{rec}}$ is rather marginal until the number of latent channels is reduced to $2^4 = 16$. We can set this as our new representation size, knowing it is sufficiently large to capture all task-critical information, and decide to address the issue of irrelevant features first.

6 Conclusion

In this paper, we introduced the notions of recognition and decision regrets, which we used to analyze the generalization performance of RL agents. Through illustrative examples, we demonstrated how these regret metrics can characterize under-specific and over-specific representations, offering a more nuanced perspective on observational overfitting in image-based RL. We highlighted how such characterization could be helpful in systematically improving the generalization performance of RL models. Future research should focus on identifying specific strategies that one can follow when an RL model fails to generalize due to over-specific vs. under-specific representations.

References

Rishabh Agarwal, Marlos C Machado, Pablo Samuel Castro, and Marc G Bellemare. Contrastive behavioral similarity embeddings for generalization in reinforcement learning. In *International Conference on Learning Representations*, 2021.

Marc Bellemare, Will Dabney, Robert Dadashi, Adrien Ali Taiga, Pablo Samuel Castro, Nicolas Le Roux, Dale Schuurmans, Tor Lattimore, and Clare Lyle. A geometric perspective on optimal representations for reinforcement learning. In *Conference on Neural Information Processing Systems*, 2019.

Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47: 253–279, 2013.

Maxime Chevalier-Boisvert, Bolun Dai, Mark Towers, Rodrigo de Lazcano, Lucas Willems, Salem Lahlou, Suman Pal, Pablo Samuel Castro, and Jordan Terry. Minigrid & Miniworld: Modular & customizable reinforcement learning environments for goal-oriented tasks. *CoRR*, abs/2306.13831, 2023.

Karl Cobbe, Oleg Klimov, Chris Hesse, Taehoon Kim, and John Schulman. Quantifying generalization in reinforcement learning. In *International Conference on Machine Learning*, pp. 1282–1289, 2019.

Will Dabney, André Barreto, Mark Rowland, Robert Dadashi, John Quan, Marc G Bellemare, and David Silver. The value-improvement path: Towards better representations for reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.

Benjamin Eysenbach, Tianjun Zhang, Sergey Levine, and Russ R Salakhutdinov. Contrastive learning as goal-conditioned reinforcement learning. In *Conference on Neural Information Processing Systems*, 2022.

Dibya Ghosh and Marc G Bellemare. Representations for stable off-policy reinforcement learning. In *International Conference on Machine Learning*, 2020.

Nan Jiang, Alex Kulesza, Satinder Singh, and Richard Lewis. The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 international conference on autonomous agents and multiagent systems*, pp. 1181–1189, 2015.

Charline Le Lan, Stephen Tu, Adam Oberman, Rishabh Agarwal, and Marc G Bellemare. On the generalization of representations in reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, 2022.

Marlos C Machado, Marc G Bellemare, Erik Talvitie, Joel Veness, Matthew Hausknecht, and Michael Bowling. Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61:523–562, 2018.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pp. 1928–1937, 2016.

Alex Nichol, Vicki Pfau, Christopher Hesse, Oleg Klimov, and John Schulman. Gotta learn fast: A new benchmark for generalization in rl. *arXiv preprint arXiv:1804.03720*, 2018.

Charles Packer, Katelyn Gao, Jernej Kos, Philipp Krähenbühl, Vladlen Koltun, and Dawn Song. Assessing generalization in deep reinforcement learning. *arXiv preprint arXiv:1810.12282*, 2018.

Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, pp. 2817–2826, 2017.

Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021.

Roberta Raileanu and Rob Fergus. Decoupling value and policy for generalization in reinforcement learning. In *International Conference on Machine Learning*, 2021.

Aravind Rajeswaran, Kendall Lowrey, Emanuel V Todorov, and Sham M Kakade. Towards generalization and simplicity in continuous control. *Neural Information Processing Systems*, 30, 2017.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Antonio Serrano-Muñoz, Dimitrios Chrysostomou, Simon Bøgh, and Nestor Arana-Arexolaleiba. skrl: Modular and flexible library for reinforcement learning. *Journal of Machine Learning Research*, 24, 2023.

Anoopkumar Sonar, Vincent Pacelli, and Anirudha Majumdar. Invariant policy optimization: Towards stronger generalization in reinforcement learning. In *Learning for Dynamics and Control*, 2021.

Xingyou Song, Yiding Jiang, Stephen Tu, Yilun Du, and Behnam Neyshabur. Observational overfitting in reinforcement learning. In *International Conference on Learning Representations*, 2020.

Adam Stooke, Kimin Lee, Pieter Abbeel, and Michael Laskin. Decoupling representation learning from reinforcement learning. In *International Conference on Machine Learning*, 2021.

Adrien Ali Taiga, Rishabh Agarwal, Jesse Farebrother, Aaron Courville, and Marc G Bellemare. Investigating multi-task pretraining and generalization in reinforcement learning. In *International Conference on Learning Representations*, 2023.

Mark Towers, Jordan K. Terry, Ariel Kwiatkowski, John U. Balis, Gianluca de Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Arjun KG, Markus Krimmel, Rodrigo Perez-Vicente, Andrea Pierré, Sander Schulhoff, Jun Jet Tai, Andrew Tan Jin Shen, and Omar G. Younis. Gymnasium, March 2023. URL https://zenodo.org/record/8127025.

Han Wang, Erfan Miahi, Martha White, Marlos C Machado, Zaheer Abbas, Raksha Kumaraswamy, Vincent Liu, and Adam White. Investigating the properties of neural network representations in reinforcement learning. *Artificial Intelligence*, 330:104100, 2024.

Amy Zhang, Nicolas Ballas, and Joelle Pineau. A dissection of overfitting and generalization in continuous reinforcement learning. *arXiv preprint arXiv:1806.07937*, 2018a.

Chiyuan Zhang, Oriol Vinyals, Remi Munos, and Samy Bengio. A study on overfitting in deep reinforcement learning. *arXiv preprint arXiv:1804.06893*, 2018b.

Xuezhou Zhang, Yuda Song, Masatoshi Uehara, Mengdi Wang, Alekh Agarwal, and Wen Sun. Efficient reinforcement learning in block MDPs: A model-free representation learning approach. In *International Conference on Machine Learning*, 2022.

Supplementary Materials

The following content was not necessarily subject to peer review.

A A Worked Example for Generalization

We now show how these definitions apply to our first worked example. The training environment $\varepsilon^{\text{train}}$ is exactly same as the original environment depicted in Figure 2. The test environment $\varepsilon^{\text{test}}$ differs in two ways. First, we modify the observation space X to a new space X^{gray} where the figures in the images are converted from red/green/blue to grayscale, as illustrated in Figure 6. Second, we start the environment from state 0 instead of state 1. The optimal decision policy π^* (i.e. the best action in each state) is the same in $\varepsilon^{\text{test}}$ as in $\varepsilon^{\text{train}}$, but due to the change in initial state the value is now $V^*_{\varepsilon^{\text{test}}} = \gamma + \gamma^3 + \ldots$, which equals 4.737 when $\gamma = 0.9$. Note that a recognition policy ρ_0 can achieve optimal performance in the training environment, i.e. $\max_\pi V_{\varepsilon^{\text{train}}}(\rho_0, \pi) = V^*_{\varepsilon^{\text{train}}}$, whenever $s_t = i$ implies $\pi'_0(x_t) = i$ for $x_t \in X$. However, different recognition policies satisfying that condition may have differing behavior in the test environment (i.e. for $x_t \in X^{\text{gray}}$).

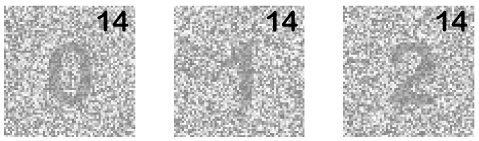

Figure 6: Modified observation space X^{gray} (at time step $t = 14$).

Example of Under-Specific Recognition Policy Suppose that $\rho'_0(\cdot)$ makes its decision based on whether there are any red, green, or blue pixels in the image, and makes a random decision otherwise. In other words, for the test environment the recognition policy does not given any useful information about the state (since the observation is grayscale). Intuitively, this is why it is deemed under-specified but we now show that it matches the definition of the term that we gave previously.

If we apply the decision policy that was optimal for $\varepsilon^{\text{test}}$ then effectively in all states we take action 0 with probability $1/3$ and action 1 with probability $2/3$ (since the representation state z_t is independent from s_t). The value in the environment for these policies is given by v_1 in the equations $v_0 = \gamma v_0(1/3) + \gamma v_1(2/3)$, $v_1 = \gamma v_0(1/3) + \gamma v_2(2/3)$, $v_2 = 1 + \gamma v_0(2/3) + \gamma v_1(1/3)$. When $\gamma = 0.9$, the solution to these equations is $(v_0, v_1, v_2) = (2.022, 2.360, 2.921)$ hence the generalization regret:

$$GR = \max_{\rho,\pi} V_{\varepsilon^{\text{test}}}(\rho, \pi) - V_{\varepsilon^{\text{test}}}(\operatorname{argmax}_\pi V_{\varepsilon^{\text{train}}}(\rho_0, \pi)) = 4.737 - 2.360 = 2.377 \tag{13}$$

Note that for this simple case, since we can calculate optimal policy in $\varepsilon^{\text{test}}$ exactly, there is no difference between GR and the empirical generalization regret $\widehat{GR}$ and so we will only focus on the former in this section.

Since the generalization regret is high, it is simple to obtain zero regret in $\varepsilon^{\text{train}}$, and the only difference between $\varepsilon^{\text{test}}$ and $\varepsilon^{\text{train}}$ is in the observation dynamics, this situation matches our definition of observational overfitting. We now decompose the generalization regret GR into its components GR^{rec} and GR^{dec}. When the recognition policy ρ_0 is used, the *optimal* decision policy for the test environment $\operatorname{argmax}_\pi V_{\varepsilon^{\text{test}}}(\rho_0, \pi)$ is simply the optimal decision policy that is not aware of the state. It is not hard to verify that this decision policy always takes action 1 and so $\max_\pi V_{\varepsilon^{\text{test}}}(\rho_0, \pi) = \gamma + \gamma^4 + \ldots$, which equals 2.989 when $\gamma = 0.9$. Hence,

$$\begin{aligned} GR^{\text{rec}} &= V_{\varepsilon^{\text{test}}}(\operatorname{argmax}_{\rho,\pi} V_{\varepsilon^{\text{test}}}(\rho, \pi)) - V_{\varepsilon^{\text{test}}}(\rho_0, \operatorname{argmax}_\pi V_{\varepsilon^{\text{test}}}(\rho_0, \pi)) \\ &= 4.737 - 2.989 \\ &= 1.748 \end{aligned} \tag{14}$$

$$\begin{aligned} GR^{\text{dec}} &= V_{\varepsilon^{\text{test}}}(\rho_0, \text{argmax}_\pi V_{\varepsilon^{\text{test}}}(\rho_0, \pi)) - V_{\varepsilon^{\text{test}}}(\rho_0, \text{argmax}_\pi V_{\varepsilon^{\text{train}}}(\rho_0, \pi)) \\ &= 2.989 - 2.360 \\ &= 0.629 \end{aligned} \tag{15}$$

Since the majority of the regret comes from GR^{rec}, this matches our definition of the recognition policy ρ_0 being *under-specified*. This is intuitive since by focusing on the colors, the recognition policy ρ_0 does not recognize the salient feature of the observation state in $\varepsilon^{\text{test}}$ (i.e. whether the number in the image is 0, 1 or 2).

Example of Over-Specific Recognition Policy We now consider a different recognition policy ρ_1 that focuses on the timestamp in the top-right corner of each image in Figure 2. In particular the representation state is defined by $\rho_1'(x_t) = 0$ if $t = 0$, otherwise $\rho_1'(x_t) = 1$ if t is odd and $\rho_1'(x_t) = 2$ if t is even. This allows us to obtain the optimal reward in $\varepsilon^{\text{train}}$ since if we apply the optimal decision policy (that takes action 1 in states $0, 1$ and action 0 in state 1), then we will follow the optimal trajectory and $s_t = \rho_1'(x_t)$ at all times t. Note that in this optimal trajectory the action sequence is $1, 1, 0, 1, 0, 1, \ldots$

However, if we move to the test environment $\varepsilon^{\text{test}}$, and apply the same recognition and decision policies, we will follow the exact same action sequence $1, 1, 0, 1, 0, 1, \ldots$ since the recognition policy is only looking at the timestamp. However, due to the change in initial state, this is not optimal in $\varepsilon^{\text{test}}$. In particular, the state sequence becomes $1, 2, 0, 0, 1, 0, 1, 0, \ldots$, i.e. we remain in a loop between states $0, 1$ rather than the states $1, 2$ where we want to be to receive the reward. The value of this sequence is $\gamma = 0.9$. Hence the generalization regret:

$$GR = \max_{\rho,\pi} V_{\varepsilon^{\text{test}}}(\rho, \pi) - V_{\varepsilon^{\text{test}}}(\text{argmax}_\pi V_{\varepsilon^{\text{train}}}(\rho_0, \pi)) = 4.737 - 0.900 = 3.837 \tag{16}$$

However, there is clearly another decision policy that still obtains the optimal reward in $\varepsilon^{\text{test}}$ even with the recognition policy ρ_1. We just need to take action 1 when $\rho_1'(x_t) = 0$, action 0 when $\rho_1'(x_t) = 1$, and action 1 when $\rho_1'(x_t) = 2$. (It can be checked that this leads to the optimal action sequence $1, 0, 1, 0, \ldots$) Hence,

$$\begin{aligned} GR^{\text{rec}} &= V_{\varepsilon^{\text{test}}}(\text{argmax}_{\rho,\pi} V_{\varepsilon^{\text{test}}}(\rho, \pi)) - V_{\varepsilon^{\text{test}}}(\rho_1, \text{argmax}_\pi V_{\varepsilon^{\text{test}}}(\rho_1, \pi)) \\ &= 4.737 - 4.737 \\ &= 0 \end{aligned} \tag{17}$$

$$\begin{aligned} GR^{\text{dec}} &= V_{\varepsilon^{\text{test}}}(\rho_1, \text{argmax}_\pi V_{\varepsilon^{\text{test}}}(\rho_1, \pi)) - V_{\varepsilon^{\text{test}}}(\rho_1, \text{argmax}_\pi V_{\varepsilon^{\text{train}}}(\rho_1, \pi)) \\ &= 4.737 - 0.900 \\ &= 3.837 \end{aligned} \tag{18}$$

In this case all of the generalization regret is concentrated on the decision regret, and so this matches our definition of the recognition policy ρ_1 being *over-specified*. This is intuitive since by focusing on the timestamp, the recognition policy ρ_1 is basing its decision on a feature that is specific to a single image, rather than the number in the image which is common to a whole image class.

B Extended Results for Maze Environments

In addition to three agent in Section 5.1, we consider two more agent with more extreme properties: (iv) *Blind* combines both *HideDoor* and *HideColors*, resulting in extremely under-specific representations. (v) *HideDoor* encodes each object-color combination as its own distinct token in a one-hot encoding scheme. Since object types, states, and colors are no longer encoded in separate channels, *OneHot* makes it much harder to learn recognition policies that ignore color, resulting in extremely over-specific representations.

An extended version of Table 1 with these two additional agents is given in Table 3. Notably, the diagonal in the similarity matrices of *Identity* is even more pronounced for *OneHot*. This

Table 3: Extended results for *Minigrid* in Section 5.2.

Agent	Pre-Filter	Performance	Representation Similarity		
			All	By Color	By Door Location
OneHot (***Extremely Over-Specific***)	[N/A]	$R = 0.002\ (0.003)$ $\widehat{GR} = 0.968\ (0.059)$ $\widehat{GR}^{\text{rec}} = 0.536\ (0.342)$ $\widehat{GR}^{\text{dec}} = 0.432\ (0.297)$	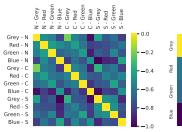	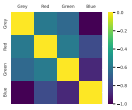	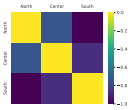
Identity (***Over-Specific***)	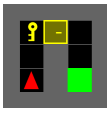	$R = 0.002\ (0.003)$ $\widehat{GR} = 0.911\ (0.171)$ $\widehat{GR}^{\text{rec}} = 0.264\ (0.388)$ $\widehat{GR}^{\text{dec}} = 0.647\ (0.365)$	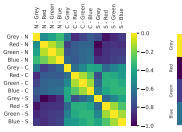	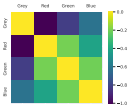	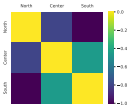
HideColors (***Ideal***)		$R = 0.000\ (0.000)$ $\widehat{GR} = 0.000\ (0.000)$ $\widehat{GR}^{\text{rec}} = 0.000\ (0.000)$ $\widehat{GR}^{\text{dec}} = 0.000\ (0.000)$	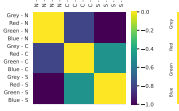	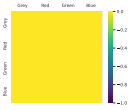	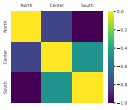
HideDoor (***Under-Specific***)		$R = 0.070\ (0.126)$ $\widehat{GR} = 0.837\ (0.157)$ $\widehat{GR}^{\text{rec}} = 0.607\ (0.227)$ $\widehat{GR}^{\text{dec}} = 0.230\ (0.217)$	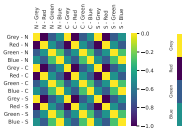	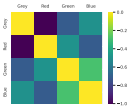	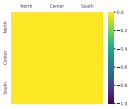
Blind (***Extremely Under-Specific***)		$R = 0.279\ (0.121)$ $\widehat{GR} = 0.279\ (0.121)$ $\widehat{GR}^{\text{rec}} = 0.256\ (0.101)$ $\widehat{GR}^{\text{dec}} = 0.023\ (0.028)$	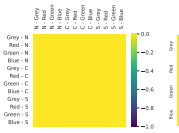	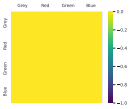	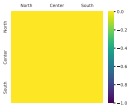

indicates that *OneHot* considers each maze in isolation without relating any features between different configurations, even more so than *Identity*. Meanwhile, *Blind* incurs a high regret even during training as it receives practically no information to base its actions on.

C Details of the Experiments

We used a host machine with a 12-core AMD Ryzen 9 7900X CPU and two Nvidia RTXA5000 (Ampere) GPUs running Red Hat Enterprise Linux 9 to spin up Ubuntu 22.04 LTS machines for training, allocating 32 GB RAM, 2 vCPUs and 1 RTXA5000 GPU per virtual machine. Using one virtual machine, the training for all experiments have taken approximately five days.

C.1 Experiments in Section 3.2

Minigrid The recognition policy ρ_{CNN} is a sequence of: a convolutional layer of channel size 16 and kernel size 2 with ReLU activations, a maximum pooling layer of kernel size 2, a convolutional layer of channel size 32 and kernel size 2 with ReLU activations, and a convolutional layer of channel size 64 and kernel size 2 with ReLU activations. The decision policy π_{MLP} is a fully-connected network with a hidden layer of size 64 and tanh activations. The PPO algorithm requires a value head in addition to the decision policy / action head; we let the value head copy the architecture of the decision policy and share the same recognition policy / encoder with the decision policy.

We consider the "`MiniGrid-SimpleCrossingS11N5-v0`" environment within the Minigrid package (Chevalier-Boisvert et al., 2023). Agents are trained in 64 parallel environments for 100,000 total time steps. We use the default hyper-parameters for the PPO algorithm in the skrl package (Serrano-Muñoz et al., 2023), except we collect 4096 total transitions in between each policy update, set the entropy loss scale as 0.01 and the value loss scale as 0.5, and activate the running standard

scalers for input images and output values. Agents are evaluated by rolling out 1024 full episodes. We repeat all our experiments five times to obtain 1-sigma error bars.

Pong The recognition policy ρ_{CNN} is a sequence of: a convolutional layer of channel size 32, kernel size 8, and stride length 4 with ReLU activations, a convolutional layer of channel size 64, kernel size 4, and stride length 2 with ReLU activations, a convolutional layer of channel size 64, kernel size 3, and stride length 1 with ReLU activations, and a fully-connected layer of size 512 with ReLU activations (known as the "NatureCNN" Mnih et al., 2015). The decision policy π_{MLP} is a fully-connected network with two hidden layers of size 64 and tanh activations. The PPO algorithm requires a value head in addition to the decision policy / action head; we let the value head copy the architecture of the decision policy and share the same recognition policy / encoder with the decision policy.

We consider the "`ALE/Pong-v5`" environment within the gymnasium package (Towers et al., 2023). Normally, the game outputs 160-by-240 full-color images. We crop and down-sample the game field to reduce these to 80-by-80 images, and also average all color channels to obtain grayscale images. Finally, we stack the four most recent images to account for memoryless agents. Agents are again trained in 64 parallel environments for $100{,}000$ total time steps. We use the default hyper-parameters for the PPO algorithm in the skrl package, except we collect 4096 total transitions in between each policy update, set the entropy loss scale as 0.01 and the value loss scale as 0.5, and activate the running standard scalers for input images and output values as well as the KL adaptive learning rate scheduler. Agents are evaluated by rolling out 256 full episodes. We repeat all our experiments five times to obtain 1-sigma error bars.

C.2 Experiments in Section 5.1

We consider the same policy architecture as in Appendix C.1 given for *Minigrid*. However this time, we consider the custom maze configurations given in Figure 4. Agents are trained in 64 parallel environments for $50{,}000$ total time steps. We use the default hyper-parameters for the PPO algorithm in the skrl package[2], except we collect 4096 total transitions in between each policy update, set the entropy loss scale as 0.01 and the value loss scale as 0.5, and activate the running standard scalers for output values. Agents are evaluated by rolling out 1024 full episodes. We repeat all our experiments five times to obtain 1-sigma error bars.

C.3 Experiments in Section 5.2

We consider the same policy architecture as in Appendix C.1 given for *Pong*. However this time, we consider a modified version of "`ALE/Pong-v5`". First, we set the repeated action probability as 0 to make the game deterministic and we reset the game as soon as the first point is scored to keep the episode lengths short. In order to obtain unseen initial states, we take between 1 and 21 null actions (inclusive, picked uniformly at random), where taking no null action corresponds to the initial state that is seen during training. While this process may still generate states that are previously seen, it guarantees generating states that are not seen as initial states before.

Images outputted by the game are processed as in Appendix C.1 to obtain 80-by-80 images. We extend these to 82-by-82 images by adding a black border at the top and at the two sides. Then, we add a frame counter at the top in the form of a progress bar: When the game is on the n-th frame, the first n pixels of the top two pixel rows are set to white (otherwise they are left black). We generate 82-by-82 images as distractions and concatenate these with the original image to obtain the final images that are then fed forward to the agents. The generating function for these distractions is as such (similar to Song et al., 2020): We start with the four-dimensional game state (consisting of the xy-coordinates of the ball and the y-coordinates of the paddles). We then pass this matrix through a neural network consisting of a fully-connected layer that brings the dimensions up to a 7-by-7

[2]The number of learning epochs during each update is 8, the number of minibatches during each learning epoch is 2, the discount factor (γ) is 0.99, the coefficient λ used for generalized advantage estimation is 0.95, the learning rate is 0.001 (an Adam optimizer is used), the clipping coefficient for the norm of the gradients is 0.5, and the clipping coefficient for computing the clipped surrogate objective is 0.2.

matrix followed by three transposed convolutions of kernel sizes $3, 4, 8$ and stride lengths $1, 2, 4$ respectively (we do not apply any activation functions, mimicking Song et al., 2020). As we have already described in Section 5.2, we randomly pick from one of two generating functions at the start of each episode. We obtain different functions to pick from through different initializations of the network of transposed convolutions. Crucially, in order to not lose the information contained within the original game state, we initialize the layers within these networks with orthogonal weights.

Finally, agents are trained in 64 parallel environments for $50{,}000$ total time steps. We use the default hyper-parameters for the PPO algorithm in the skrl package, except we collect 4096 total transitions in between each policy update, set the entropy loss scale as 0.01 and the value loss scale as 0.5, and activate the running standard scalers for input images and output values as well as the KL adaptive learning rate scheduler. Agents are evaluated by rolling out 256 full episodes. We repeat all our experiments five times to obtain 1-sigma error bars.

C.4 Experiments in Section 5.3

We consider the same setup as in Section 5.1. Only difference is that we vary the channel size of the last convolutional layer. During these experiments, especially for models with very small channel sizes, PPO tends to output unstable solutions. To account for the increased variation, we repeat all our experiments ten times instead of five to obtain 1-sigma error bars.