

LOSS OF PLASTICITY IN CONTINUAL DEEP REINFORCEMENT LEARNING

Zaheer Abbas^{*1}, Rosie Zhao², Joseph Modayil¹, Adam White^{3,4,5}, Marlos C. Machado^{3,4,5}

¹ DeepMind ² Harvard University ³ University of Alberta, Department of Computing Science

⁴ Alberta Machine Intelligence Institute (Amii) ⁵ CIFAR AI Chair

ABSTRACT

In this paper, we characterize the behavior of canonical value-based deep reinforcement learning (RL) approaches under varying degrees of non-stationarity. In particular, we demonstrate that deep RL agents lose their ability to learn good policies when they cycle through a sequence of Atari 2600 games. This phenomenon is alluded to in prior work under various guises—e.g., loss of plasticity, implicit under-parameterization, primacy bias, and capacity loss. We investigate this phenomenon closely at scale and analyze how the weights, gradients, and activations change over time in several experiments with varying experimental conditions (e.g., similarity between games, number of games, number of frames per game). Our analysis shows that the activation footprint of the network becomes sparser, contributing to the diminishing gradients. We investigate a remarkably simple mitigation strategy—Concatenated ReLUs (CReLUs) activation function—and demonstrate its effectiveness in facilitating continual learning in a changing environment.

1 INTRODUCTION

Deep Reinforcement learning (RL) approaches have been successfully used in several real-world applications, from discovering superior algorithms for video compression (Mandhane et al., 2022) and matrix multiplication (Fawzi et al., 2022) to designing superior floorplans for Google’s tensor processing units (Mirhoseini et al., 2021); from efficiently controlling cooling systems (Luo et al., 2022) and thermal power generators (Zhan et al., 2022) and managing inventories (Madeka et al., 2022) to navigating balloons in the stratosphere (Bellemare et al., 2020) and automatically producing plasma configurations for Tokamak nuclear fusion reactors (Degrave et al., 2022).

Largely, these systems can be characterized as a search for a fixed policy that once deployed no longer changes. The search is typically done offline using a batch of historical data or more often on a simulator (e.g., Madeka et al., 2022; Zhan et al., 2022; Mandhane et al., 2022; Fawzi et al., 2022; Degrave et al., 2022; Bellemare et al., 2020). Indeed, in some cases the pursuit of a stationary policy is wholly appropriate as the policy does not need to change once deployed (e.g., a policy that multiplies matrices with the desired computational complexity).

There are many applications, however, in which a stationary policy can quickly become ineffective once the system is deployed and more data become available. In cooling-system controllers, for instance, change in weather patterns and sensor drift requires continual learning to produce effective policies (Luo et al., 2022). More generally, continual adaptation is unavoidable in situations in which the world around the learning system continually changes (Sutton et al., 2007).

Most modern deep RL agents are based on ideas and algorithms from supervised learning which can be problematic in the continual learning settings. In supervised learning, data is sampled independent and identically distributed (IID) from a fixed data set (Hadsell et al., 2020). Most deep RL agents make use of large replay buffers to help approximate the stationary supervised learning regime, even though RL problems are inherently non-stationary as the data distribution shifts with changing policies. Reinforcement learning agents that are embedded in a complex world must deal with non-stationarities that manifest over timescales much longer than what can feasibly be stored in a replay buffer (Lesort et al., 2020; Luo et al., 2022), emphasizing the need for continual learning solutions that can work across timescales.

In supervised learning, recent work has shown that neural networks learn poorly when data become available incrementally (Ash & Adams, 2020), and neural networks lose their ability to learn altogether when confronted with distribution shifts over time—a phenomenon characterized as *loss of plasticity* (Dohare et al., 2021). In a setting similar to Ash &

* Corresponding author: zaheersm@google.com.

Adams’s (2020), Berariu et al. (2021) investigated different aspects of the learning problem that can impact a neural network’s plasticity, showing how distribution shift and different learning phases can degrade performance. In RL, neural networks lose capacity over time as they estimate continually changing value functions resulting from changing policies (Kumar et al., 2021)—possibly related to changes in the curvature of the loss landscape (Lyle et al., 2022). Others have shown that deep reinforcement learning agents *overfit* to experiences gathered early in training, being unable to adjust as more experience becomes available—a phenomenon known as *the primacy bias* (Nikishin et al., 2022).

In this paper, we contribute to the growing literature on catastrophic loss of plasticity in several ways. First, we demonstrate catastrophic loss of plasticity in a performant deep reinforcement learning system, Rainbow (Hessel et al., 2018), on a variation of the Arcade Learning Environment (ALE) (Bellemare et al., 2013; Machado et al., 2018) that involves non-stationarities from changing Atari 2600 games and game modes over time. Some of the experiments presented in this paper are run up to two billion environment interactions spanning 50 days of wall clock time. The result depicted in Figure 1 provides an example of the larger set of performance results presented in this paper. Next, we perform an extensive analysis of weights, gradients, and activations sampled over the lifetime of the experiment. The analysis augments the existing findings in the literature and points to a remarkably simple mitigation. Finally, we evaluate one mitigation—the concatenated ReLU activation—and highlight its effectiveness.

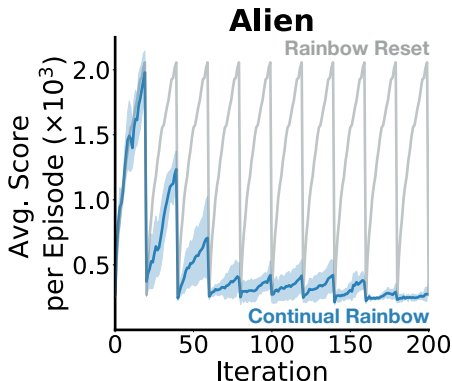


Figure 1: **Continual learning and the loss of plasticity in an Atari game:** online performance of a Rainbow agent (in blue) within a single game (Alien) when repeatedly playing a sequence of games. Starting with Alien, the agent plays a sequence of 10 games, training on each game for 20M consecutive frames, and revisits Alien after completing the cycle. Rainbow’s performance deteriorates on successive visits of the game, indicating a loss in its ability to learn. For calibration, the gray line depicts the performance of an agent if it were reset after each switch, learning from scratch for 20M frames on each visit to the game. The reset agent’s performance would start low and reach the same level each time.

2 PRELIMINARIES

We study agents that continually interact with their environment and learn to achieve a goal. In particular, the agent’s objective is to select actions based on some state, $S_t \in \mathcal{S}$, to maximize a scalar reward signal, $R_t \in \mathbb{R}$, over time. The agent’s action, A_t , is determined by its policy, $\pi(A_t|S_t)$, which is adapted towards the goal of maximizing total discounted future reward: $G_t \doteq \sum_{i=0}^{\infty} \gamma^i R_{t+1+i}$, where $\gamma \in [0, 1)$.

In this paper we focus exclusively on methods that learn value functions to obtain the policy: Deep Q-learning (DQN) and related variants. Put simply, a DQN agent continually adapts a parameteric estimate of future reward called the *value function*: $\hat{q}_{\mathbf{w}}(s, a) \approx \mathbb{E}[G_t|S_t = s, A_t = a]$, where \mathbf{w} are the parameters of a neural network. On each timestep, DQN (1) selects an action according to some exploratory policy (e.g., ϵ -greedy), (2) stores the most recent experience in replay buffer, and (3) samples a mini-batch from the replay buffer, and updates the weights with gradients computed with the backpropagation algorithm: $\nabla_{\mathbf{w}} \mathbb{E}[(R_{t+1} - \gamma \max_{a \in \mathcal{A}} \hat{q}_{\bar{\mathbf{w}}}(S_{t+1}, a)) - \hat{q}_{\mathbf{w}}(S_t, A_t)]^2$, where $\bar{\mathbf{w}}$ are the parameters of the target network that are periodically set equal to \mathbf{w} . Often (e.g., in Atari 2600 games) the agent does not observe S_t , but instead an incomplete summary of the state such as the four most recent game frames.

The majority of our experiments focus on the Rainbow agent: an implementation based on DQN (Mnih et al., 2015), which is highly tuned to play Atari 2600 games (Hessel et al., 2018). In particular, Rainbow combines several algorithmic components including distributional RL (Bellemare et al., 2017), noisy networks for exploration (Fortunato et al., 2018), prioritized experience replay (Schaul et al., 2016), dueling networks (Wang et al., 2016), n-step returns (Sutton & Barto, 2018), and a few other ideas. The performance of Rainbow is representative of the best performing model-free, synchronous (single environment) RL implementations available. Because of dueling networks, Rainbow’s neural network has two components, one that estimates the value function, $\hat{v}_{\mathbf{w}_1}(s) \approx \mathbb{E}[G_t|S_t = s]$, and one that estimates the advantage, $\hat{a}_{\mathbf{w}_2}(s, a)$, such that $\hat{v}_{\mathbf{w}_1}(s) + \hat{a}_{\mathbf{w}_2}(s, a) \approx \mathbb{E}[G_t|S_t = s, A_t = a]$. In this paper, we term

the shared network layers, before the split for the two different components, *convolution network*, the set of layers used in \hat{v}_{w_1} 's part *value network*, and those used in \hat{d}_{w_2} 's part *advantage network*.

In this paper we focus our empirical analysis on a subset of Atari 2600 games using the Arcade Learning Environment (ALE) (Bellemare et al., 2013; Machado et al., 2018). The ALE is a standard domain for evaluating deep reinforcement learning algorithms and the two specific algorithms we study, Rainbow and DQN, were originally developed for Atari 2600 games. The traditional evaluation protocol in the ALE treats all games independently: the agent interact with each Atari 2600 game for 200 million frames, reporting the final performance achieved by the agent. We abide by Machado et al.'s (2018) recommendations of evaluation protocol, including the use of stochasticity, ignoring the lives signal, and reporting the average performance during training. We diverge from the standard practice: a single agent is trained to play a sequence of games—one after the other—with a single neural network (which we describe in detail in the next section). In most prior work agents are forced to follow a specific schedule to ensure samples of all games are seen often during training, such as ensuring each batch contains samples drawn uniformly from all the available games (e.g., Reed et al., 2022). The setup explored in this paper is designed to mimic non-stationarity in continual learning tasks.

3 DEMONSTRATING LOSS OF PLASTICITY

In this section we describe Switching ALE, a continual variant of the ALE, and we use it to evaluate the performance of Rainbow under varying degrees of non-stationarity, demonstrating how the Rainbow agent loses plasticity over time.

3.1 ADAPTING THE ALE FOR CONTINUAL LEARNING

Our objective is to evaluate how well deep RL methods can continually adapt in a changing world. We consider a simple experimental regime that is derived from the popular ALE benchmark. The common practice in ALE is to develop performant algorithms that learn to play each game separately in a fixed number of interactions—typically 200M frames. Here, following the continual learning experimental practice of using a collection of tasks (Zenke et al., 2017; Thrun, 1995; Taylor & Stone, 2009; Yu et al., 2020; Khetarpal et al., 2022), the agent learns on each game for a fixed number of frames and then switches to the next game *without* resetting the learned weights nor the replay buffer contents in between.

Specifically, the environment cycles through a fixed sequence of games, (e.g. Alien, Atlantis Boxing, ..., Alien, ...) for a fixed number of agent-environment interactions between any two switches (specifically 20M frames)¹. Consequently, each game is repeated numerous times during the course of the experiment. We call the agent's experience playing a single game a visit. For instance, given a budget of 1 billion frames and a set of 10 games and 20M frames per game, it will take 200M steps to cycle through all games. The agent will get to play (visit) each game 5 times over the length of the experiment, spending 180M frames between successive plays of the same game. Note, the duration of our experiments far exceeds the replay capacity of 1M frames, typical in ALE experiments. We refer to this experimental set up as Switching ALE, or S-ALE for short.

Naturally, there have also been other proposals for continual RL tasks, such as Jelly Bean World (Platanios et al., 2020) and Continual World (Wolczyk et al., 2021), which differ from ours in several dimensions such as the scale in which experiments are performed, the dimensionality of the observation space, the diversity of tasks, and the actual protocol used to induce non-stationarity.

3.2 LEARNING PERFORMANCE IN S-ALE

Figure 2 depicts the performance of a Rainbow agent in S-ALE when cycling through 5 games, spending 20M frames per visit. Prior work shows that Rainbow achieves impressive scores when trained on individual Atari 2600 games (c.f. Hessel et al., 2018). In the S-ALE setup, on the other hand, the agent performs relatively well when playing a game for the first time, but the performance dramatically decreases in most games on successive visits. We report results for each game in separate plots because the rewards in different games are not at the same scale. Moreover, the learning curve for each game (each subplot) reports only the performance of the agent while playing that particular game. The result for Alien, for example, only shows online learning performance when the agent is playing Alien; the agent's

¹When cycling through 5 games, the order in which these games are seen is: Alien, Atlantis, Boxing, Breakout, and Centipede. When cycling through 10 games, the order in which these games are seen is: Alien, Atlantis, Boxing, Breakout, Centipede, Double Dunk, Freeway, Pong, Space Invaders, and Tennis. In general, the order in which the games are cycled through is the same in which we present the results.

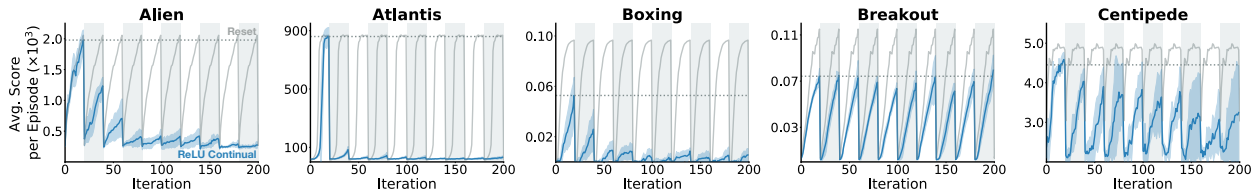


Figure 2: **Rainbow’s loss of plasticity on a repeating sequence of 5 Atari 2600 games.** Each subplot reports learning within a single game. The blue curves show Rainbow learning as it cycles through 5 games in a fixed sequence. The agent interacts with each game for 20M frames in every iteration of the cycle; the cycle restarts every 100M frames. The gray line depicts the performance achieved by an idealized reset (described in text) that represents the performance of an agent that is reset every 20M frames; as if there would be no impact from other games or previous visits to the current game. The dotted line marks the average performance achieved by the end of the first visit—a visual reference to highlight performance drop over successive visits.

performance in other games is not depicted in the same subplot. Finally, each curve is averaged over five runs and the shading depicts their standard deviation.² The results show that every time the agent visits a game again the agent learns more slowly than before, and the policy at the end of visit is typically worse, ultimately collapsing to showing no signs of improvement.

To help make it easier to contextualize these results, consider an agent that completely forgets how to play a game after each switch, but maintains indefinitely its ability to learn the current game it is visiting—relearning anew from scratch. Furthermore, let us assume that the experience the agent has playing a game has no impact over its policy when learning others game (no interference, no generalization). We call this the *reset agent* as it equivalent to resetting the network weights, replay buffer, and optimizer state after every switch such that it is always learning from scratch. This agent is far from ideal as it is forgetting absolutely everything it has experienced before, but because of its never-ending learning potential, its performance will always recover in expectation to the same level seen at the end of the last visit. Ideally any agent able to continually learn in the environment should be able to match and outperform the reset agent. But the results in Figure 2 show that Rainbow performs worse than the reset agent over time.

The results we report for the reset agent are idealized. For every game, we record the performance of a freshly initialized Rainbow agent on that game and we report this performance as the agent’s first-visit performance for every other successive visit. Consequently, all the bumps of the gray curve in Figure 2 are identical. Because Alien is always the first game in the sequence of games the agent plays, the blue and gray curves overlap only for this game. In the other games, the agent’s performance is already impacted by previous learning. In fact, some loss of plasticity could potentially already have occurred from visits to previous games. We use such a reset agent throughout the paper.

We hypothesize that Rainbow is suffering from both catastrophic interference, the phenomenon where prior learning overrides previous learning (French, 1999; Liu et al., 2020), and loss of plasticity, the phenomenon where the network loses the ability to learn. On revisiting a game, we see that the agent not only needs to start relearning, it typically achieves much lower performance compared with the previous visits. Importantly, we observe this phenomenon consistently in different settings, as we vary the number of games the agent cycles through, or the frequency in which game switches happen. Figure 10 (in the Appendix) shows a similar pattern when the agent cycles through 10 games. We also observe performance degradation when we change the number of frames per visit. We see similar, if less pronounced results (as expected), in the Appendix in Figure 11, with 10M frames per visit—thus many more visits to each game over 1B frames. With 50M frames per visit in Figure 12, also in the Appendix—only two visits to each game over the experiment—we see similar degradation compared to the our first result with 20M frames per visit.

Loss of plasticity is not unique to Rainbow, when cycling through 10 games with 20M and 50M frames per visit, we also observe the same performance degradation with the DQN, as shown in Figures 13 and 14 in the Appendix.

3.3 VARYING THE DEGREE OF NON-STATIONARITY: LOOPING THROUGH GAME MODES

Here we consider a milder form of non-stationarity, induced by changing modes within a single game. In the previous experiment we cycled through a sequence of games. The degree of non-stationarity due to changing games is presumably dramatic—the distribution of observations, the transition dynamics, even the scale of reward changes across

²Ideally, more runs would be used, but as mentioned above, these experiments took over a month to run. Consider these results demonstrative in nature: showing how deep RL agents can fail in a continual learning setup. Our results cannot provide concrete evidence that such failures will always occur.

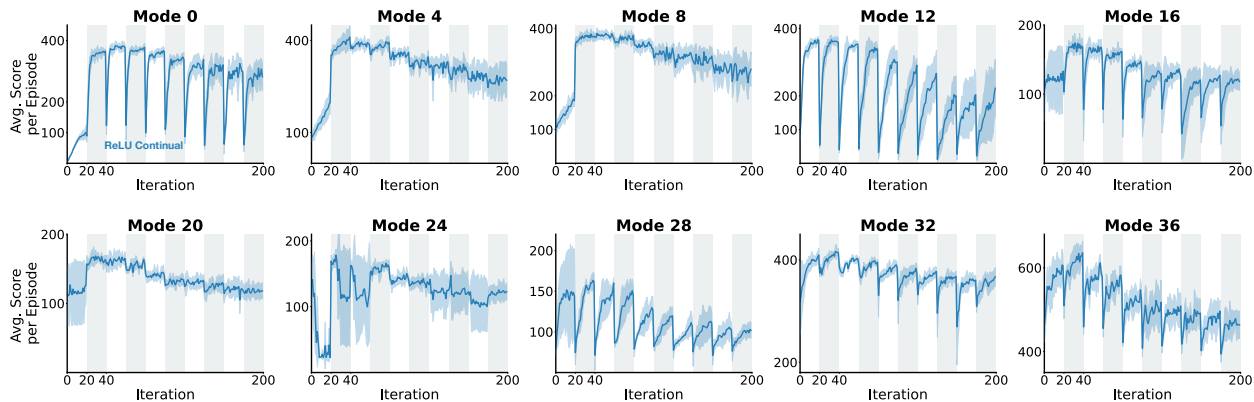


Figure 3: **Loss of plasticity in Breakout on a sequence of 10 game modes.** The structure of this plot is similar to Figure 2. The blue curves show Rainbow learning as it cycles through 10 different game modes of Breakout in a fixed sequence. The agent interacts with each game mode for 20M frames in every iteration of the cycle; the cycle restarts every 200M frames. We again observe loss of plasticity in this sequence of modes despite the game modes being similar to each other: similar dynamics and observation space (see the main text for details).

games (c.f. Machado et al., 2018; Farebrother et al., 2018). There is some evidence, however, of successful transfer of policies across Atari 2600 game modes (Rusu et al., 2022; Wang et al., 2022), hinting at a high degree of similarity.

We experimented with changing game modes in 3 Atari 2600 games (Breakout, Freeway, and Space Invaders) and, interestingly, we found clear evidence of loss of plasticity in only one game. In Breakout, Figure 3, the performance deteriorates over successive visits to the same game mode, paralleling our results in changing games. Looking closely at Figure 3, the most dramatic performance loss is observed in game modes 4, 8, 20, and 36. In Breakout, the agent must learn to reflect the ball off a paddle which it can only move left or right along the bottom of the screen. In game mode 8, the agent has two additional actions: catch and release the ball, which affords significantly more control. Similarly, modes 4, 30, and 36 allow the agent to steer the ball after reflection off the paddle. These four modes present the biggest changes in dynamics, potentially explaining why the performance in these modes matches our first result (cycling through different games).

In Freeway (Figure 15), the agent recovers faster when revisiting a game mode and exceeds performance of prior visits. Finally, in Space Invaders (Figure 16), the agent continues to perform well across visits to the same game mode, except in one game mode. All figures referenced in this paragraph are in the Appendix. These results suggest that loss of plasticity is somewhat related to the similarity between the different tasks an agent might be faced with, but it is difficult to characterize such similarities beforehand. The different game modes induce, one way or another, a change in dynamics, but specific ones (as in Breakout) seem to have a bigger impact than others.

4 CHARACTERIZING LOSS OF PLASTICITY

In this section we take a closer look at how the internal components of Rainbow evolve over time. Before we start, it is helpful to summarize our findings. As shown earlier, the performance progressively worsens each time the Rainbow agent revisits a game and, in some cases, it collapses altogether.

We observe the first hint of reduced plasticity by inspecting how much the network weights change each time the agent visits a game. In the game Alien, for example, the *weight change* consistently shrinks with additional visits (c.f. Figure 4a). More critically, the weight change continues to diminish even though the (distributional) loss remains large (Figure 4b). One would expect a large loss to result in large gradients, but we find that the magnitude of gradient also diminishes on every successive visit (Figure 4c). Finally, looking at the activations over time, we see fewer and fewer units are active (Figure 4d). This makes sense. The Rainbow agent uses Rectified Linear Units (ReLU). If a ReLU is not active (i.e. zero output value), then it does not contribute to the change in its incoming weights (via the chain rule). A learning system that is made up of unchanging weights, irrespective of a large global loss, can be well thought of as having lost its ability to learn new things—the network is losing plasticity.

Next we discuss how we computed each of the statistics discussed above and present more detailed comparisons.

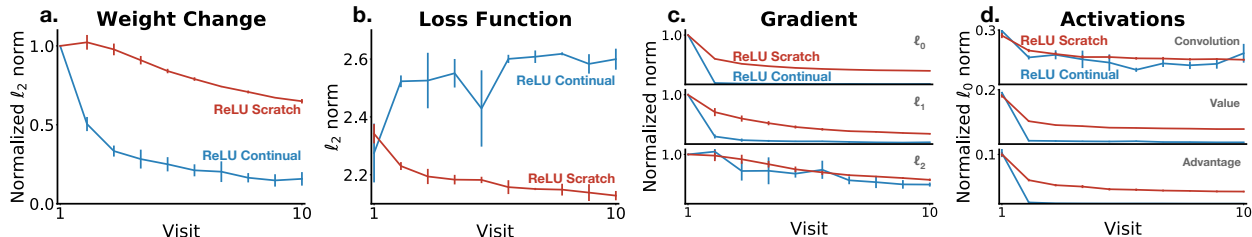


Figure 4: **Weight Change in Alien.** Normalized weight change of the Rainbow network over successive plays of Alien. **Loss Function.** Loss at halfway point over successive plays. For calibration, we include a scratch baseline that plays Alien exclusively and uninterrupted and sees the same number of game frames (see text for details). **Gradient.** Normalized gradient norms (ℓ_2 , ℓ_1 , and ℓ_0) aggregated across layers. **Activations.** ℓ_0 norm of the activations.

4.1 WEIGHT CHANGE DIMINISHES OVER TIME

We compare the weight change over every visit in S-ALE to the weight change produced by an agent learning from scratch in a conventional (non-continual) ALE experiment. Specifically, we compute the ℓ_2 -norm of the weight change for every visit from when the agent start playing Alien until it is halfway through the current visit (i.e. over 10M frames). We compute the weight change for the *scratch agent* by simply stacking visits one after the other and no other game is played in between. Results are shown in Figure 4a.

We normalize the weight changes to make this comparison meaningful, because Rainbow uses a multi-layered network and the weight magnitudes across layers are not scale invariant. We use the weight change observed in the first visit to scale the norms of weight change in subsequent visits (thus the scaled norm can be greater than one). We aggregate the result from all layers with a weighted average, weighted by the number of weights in that layer.

We contrast weight change with how the loss changes during training. We compute the loss depicted in Figure 4b by averaging the loss halfway through each visit: halfway through each visit, we average the empirical loss on 100 mini-batches. In the continual set up, the weights change diminishes sharply—on the tenth visit, the weight change is 20% of that in the first visit (blue curve in Figure 4a). This is notable because the loss, on the other hand, grows large. In contrast, the scratch agent’s weights (the red curve in Figure 4a) continue to change much: the weight change on the tenth visit is still 75% of that in the first visit (consistent with the decreasing loss for the scratch agent).

4.2 GRADIENT COLLAPSE

The norm of the gradients helps us understand why the weights stop changing, but the loss remains high. We summarize the gradients as follows: at the halfway point of each visit, we compute the average layer-wise norm of the gradients for the next 100 updates. We employ the same normalization and aggregation approach as the weight change result. Figure 4c reports the ℓ_0 , ℓ_1 , and ℓ_2 norms of the gradients for the continual and scratch agents. The ℓ_0 and ℓ_1 norms of the gradients for the continual set up (blue curves) decay much faster than the corresponding norms of the scratch agent (red curves). The gradients inside Rainbow’s network collapse almost to zero in S-ALE! Note the decay in ℓ_2 is not as drastic, likely because ℓ_2 norm’s susceptibility to outliers due to squaring.

We report the value of the gradients in Figure 4c, before they are scaled by the optimizer. However, note that Figure 4a shows the weights also stop changing, thus the scaling introduced by the optimizer is not making a difference. This is expected as Adam tracks the exponentially decaying average of the gradient and its square. If the gradients collapse to zero (as shown), the decaying average of the gradient and its square also collapses to zero.

4.3 ACTIVATION COLLAPSE

Rainbow exhibits poor performance, increasing training loss, reducing weight change, and dying gradients—inspecting the activations in the network shed some light on the collapse of gradients and it also points toward a solution. At halfway point of each visit, we record the internal activations of the network for the same 100 mini-batches that were used for the gradient inspection. The network in the Rainbow agent consists of a stack of convolution layers whose output is fed into a value network and an advantage network. The output of the convolution stack consists of 3136 units; the value and the advantage network have one hidden layer each of size 512. In Figure 4d, we show the ℓ_0 norm of each layer’s activations scaled by the number of units, average across 100 mini-batches of size

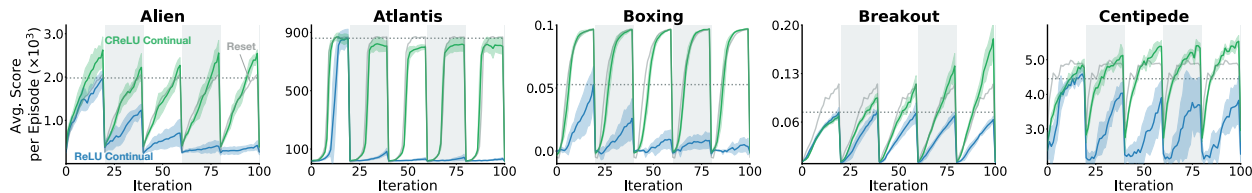


Figure 5: **CReLU activation function mitigates Rainbow’s loss of plasticity in S-ALE.** The structure of this plot is identical to Figure 2. The blue curves show the performance of the Rainbow agent with ReLU activations in the continual learning setting as it cycles through 5 different Atari 2600 games in a fixed sequence. The agent interacts with each game for 20M frames in every iteration of the cycle; the cycle restarts every 100M frames. The green curve shows the performance of Rainbow with CReLU activations in the S-ALE experimental setup. CReLUs typically do not result in faster initial learning (when compared to the reset agent) but maintains the ability to learn, continuing to match (or exceed, as in Boxing) the performance of reset agent and its own performance on prior visits—a basic requirement for continual learning.

32. We see that the ℓ_0 norm of activations in the advantage and value network collapses: a very small fraction of the units of the Rainbow’s network in the continual setting produce non-zero value. If a ReLU is zero, then it does not contribute to the change in its incoming weights, hindering adaptation. Though activations of the convolution stack do not collapse, the adaptation of convolution layers relies on the gradients to flow back from the advantage and value layers—consequently, task changes in S-ALE will impact the convolutional layers less and less over time.

5 MITIGATING LOSS OF PLASTICITY WITH CONCATENATED RELUS

In this section we explore a simple approach to increase and maintain network activations and thus mitigate the loss of plasticity. Recall that the internal activations of Rainbow’s network decayed to the point where only a small fraction of the units produced non-zero values (less than 1% in fact). This collapse in activation resulted in diminishing gradients and prevented further change to the weights, thus, impeding learning. Here we consider a simple change to the network architecture that can prevent activation collapse: Concatenated Rectified Linear Units (CReLUs). Originally proposed to understand and improve convolutional architectures for object recognition (Shang et al., 2016), CReLU concatenates the input with its negation and applies ReLU to the result: $\text{CReLU}(x) \doteq [\text{ReLU}(x), \text{ReLU}(-x)]$. CReLU outputs twice the number of input signals and ensures that, for every input signal, one of the two output signals is nonzero, except when the input signal is exactly zero.

Evaluating CReLUs requires us to pay attention to the effective capacity of the network. In a given layer, replacing ReLU with CReLU doubles the number of activations while keeping the number of parameters the same. This doubling can potentially double the number of parameters in the following layer. To isolate the contribution of CReLU in mitigating the loss of plasticity, we experiment with two approaches. The first approach controls for the number of input signals before the application of the activation function (*invariant input dimension*). For the CReLU variant of continual Rainbow, this doubles the number of parameters in every layer (except the first one). The second approach controls for the number of output signals after the application of the activation function (*invariant output dimension*). This halves the number of parameters in every layer (except the last one) in the CReLU variant. We evaluate both architectures in our S-ALE setup with no change in hyper-parameters.

The performance of the Rainbow agent with CReLU and *invariant input dimension* is depicted in Figure 5. The results show that Rainbow-CReLU—unlike Rainbow-ReLU—continues to maintain or improve its performance on successive visits in all five games. On every repeated visit to a game, having forgotten what was learned on the previous visit, Rainbow-CReLU (like Rainbow with ReLUs) must learn to play the game again. The Rainbow-CReLU agent demonstrates improved plasticity by matching (or exceeding, in Breakout and Centipede, for example) the performance of the reset agent and its own performance in prior visits. The results for the *invariant output dimension* can be seen in Figure 17 in the Appendix. This variant of Rainbow-CReLU with less parameters also clearly demonstrates maintenance of plasticity.

It would be reasonable to suspect that the improvement in plasticity is simply due to the general superiority of CReLU over ReLU, but this does not seem to be the case. Figure 6 shows that ReLU-Rainbow (i.e. ReLU Scratch) performs comparably in the conventional (non-continual) ALE. For completeness, we also compared continual Rainbow-CReLU against the scratch Rainbow-CReLU agent (as opposed to the scratch Rainbow-ReLU in the plots above) to

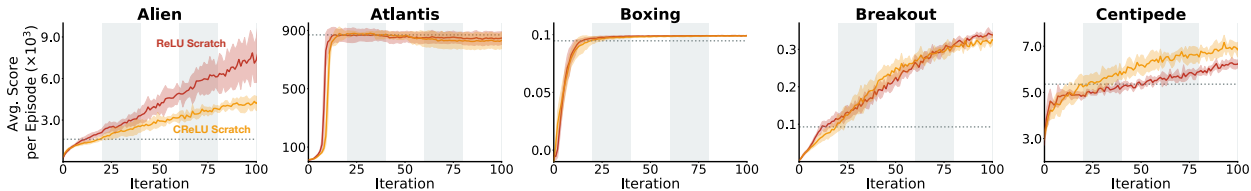


Figure 6: **Comparing Rainbow with CReLUs and Rainbow with ReLUs.** When training a Rainbow agent from scratch in a single game for 100M frames, CReLU and ReLU activations lead to comparable performance.

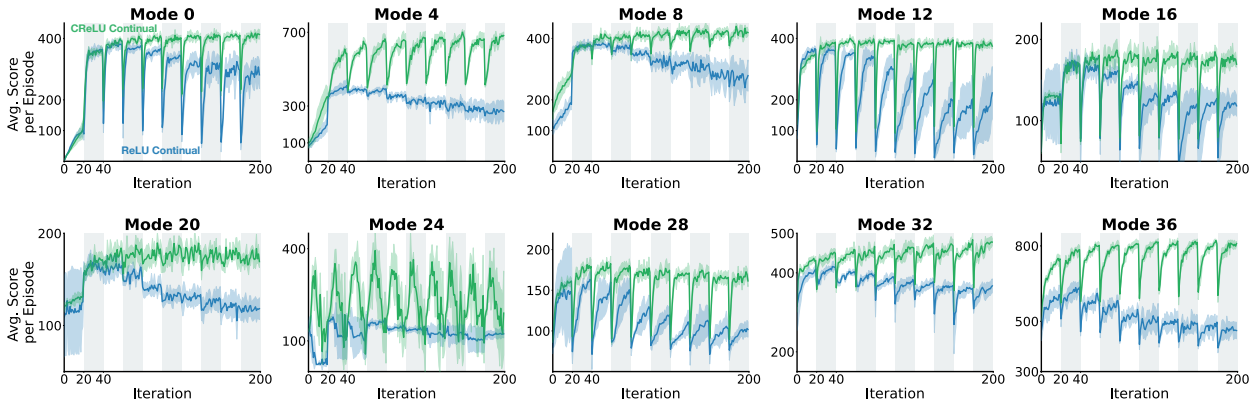


Figure 7: **Mitigating loss of plasticity in Breakout game modes with CReLU activations.** The structure of this plot is identical to Figure 3. The blue curves show the performance of the Rainbow agent with ReLU activations in the continual learning setting as it cycles through 10 different game modes of Breakout in a fixed sequence. The agent interacts with each game mode for 20M frames in every iteration of the cycle; the cycle restarts every 200M frames. The green curves show the performance of Rainbow with CReLU activations in the same continual learning setting.

quantify the impact in performance that cycling through multiple games instead of learning only from scratch has when using CReLU activations. The result can be found in Figure 18 in the Appendix.

We see a similar pattern when we evaluate Rainbow-CReLU for sequential training on Breakout game modes. As depicted in Figure 7, Rainbow-CReLU continues to maintain or improve its performance on successive visits in all 10 game modes. Rainbow-CReLU shows no significant improvement over Rainbow across the game modes of Space Invaders, presumably because Rainbow did not experience significant loss of plasticity (see Figure 19 in the Appendix for the evaluation of Rainbow-CReLU in the different game modes of Space Invaders).

Finally, we inspect the internal statistics of the Rainbow-CReLU agent to better understand the impacts of CReLU on network dynamics. They constitute clear evidence of improvement in plasticity. As before we focus on the game Alien. First, Figure 8d shows that about half of the activations continue to produce non-zero values; CReLU is preventing activation collapse (the green curves). It does so by construction: for every input signal, one of the two output signals is nonzero, except when the input signal is exactly zero (unlikely because of the 32-bit floating point representation of real numbers).

Figure 8c shows that the relative norm of the gradients stays large (as we would expect if activations did not collapse). The ℓ_0 norm does not diminish relative to the first visit. Relative ℓ_1 norm does reduce, however, it does not collapse like with ReLUs, reducing much more slowly than Scratch; and the rate of decay becomes smaller over visits. Figure 8a shows that weights in Rainbow-CReLU continue to change significantly relative to the first visit all the way through ten visits: as we would expect if gradients did not collapse. The relative weight change for Rainbow-CReLU is larger than not only Rainbow-ReLU but also the Scratch baseline, which makes sense given that the loss function observed with CReLUs is larger than the loss function observed with ReLU Scratch. Figure 8b shows the sampled loss of Rainbow-CReLU. It does not grow to be as large as Rainbow-ReLU, likely due to its improved plasticity that allows it to continue to learn. Note, as aforementioned, that the training loss of continual Rainbow with CReLUs is still higher than Scratch; this is the topic of the next section.

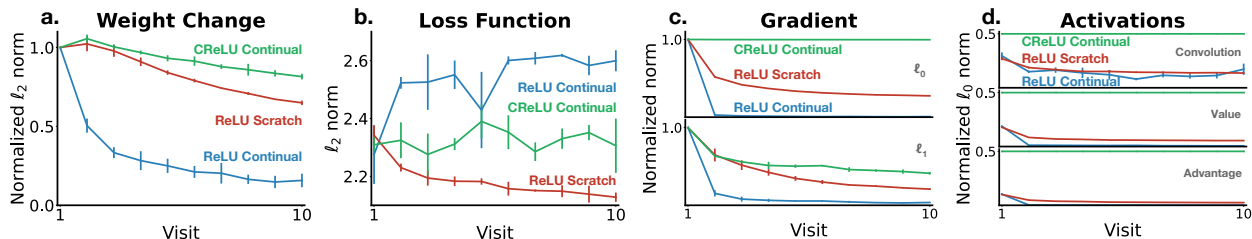


Figure 8: **Weight Change in Alien.** Normalized weight change of the Rainbow network over successive plays of Alien. **Loss Function.** Loss at halfway point over successive plays. For calibration, we include a scratch baseline that plays Alien exclusively and uninterrupted and sees the same number of game frames (see text for details). **Gradient.** Normalized gradient norms (l_2 , l_1 , and l_0) aggregated across layers. **Activations.** l_0 norm of the activations.

6 CATASTROPHIC FORGETTING: AN UNRESOLVED CHALLENGE IN CONTINUAL LEARNING

Up to this point we have discussed a large number of results, and thus it is useful to summarize what we have learned. We have focused on catastrophic loss of plasticity as the core phenomenon, demonstrating it in the S-ALE setup. We analyzed the learning dynamics over time and evaluated a simple mitigation strategy. In this section, we discuss how the same S-ALE experiments relate to the well-established phenomenon of catastrophic forgetting (French, 1999).

In all the results discussed in this paper, every time the Rainbow agent revisits a game (or mode), it shows no retention of what was previously learned: the agent forgets. The CReLU mitigation does not address this issue and we see no reason to expect as much. Consider the results in Figure 9. The Rainbow agent with CReLU activations, when cycling through five games, can improve the performance over successive visits in some cases (green curve), but the improvement appears modest when we compare with an agent that learns uninterrupted on a single game with the same amount of experience (red curve). Figure 20 in the Appendix shows similar results for a ten-game sequence and Figure 18, also in the Appendix, presents similar results when using CReLU Scratch as a baseline.

It is unclear if it is realistic to expect agents such as Rainbow—designed for non-continual ALE—to remember what they learned over 100M frames ago and perform transfer learning between visits. Optimistically, we might hope that an agent that trains on many games might learn a better representation and a better policy than an agent that only trains on a single game—even if the total number of frames were the same. Unfortunately, our results show there is a widening gulf between the idealized agent that learns uninterrupted (the red curve in Figure 9) and continual Rainbow. The path forward seems clear: synthesizing continual learning approaches that address stability-plasticity dilemma: addressing both catastrophic forgetting and the loss of plasticity.

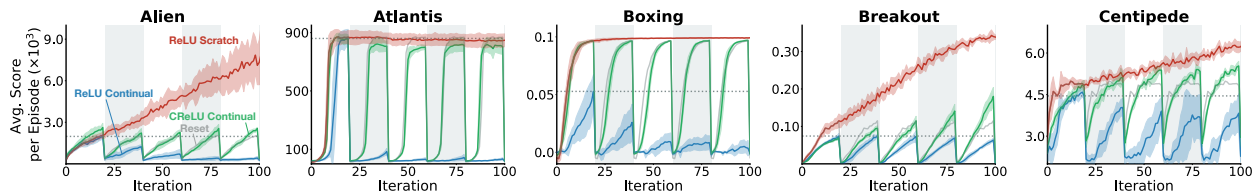


Figure 9: **Catastrophic forgetting and the loss of plasticity in S-ALE.** The structure of this plot is identical to Figure 5. Here we also include in red the performance of a conventional Rainbow agent (with ReLU activations) training on each game, individually and uninterrupted, for 200M frames—the conventional non-continual ALE setup. We expect no notable loss of plasticity in this setting, and thus the red curve sets a clear performance bar for the other agents. This provides clear context for the CReLU result: for performance within a single game, there is no benefit from learning the other games, and barring the two games in which good performance can be reached within a single visit (Atlantis and Boxing), the gulf between the performance of scratch baseline and continual CReLU gets wider over successive visits—a sign of catastrophic interference.

7 CONCLUSIONS AND FUTURE WORK

Continual adaptation is essential as no amount of prior learning is sufficient in complex and continually changing environments. Most of the solution methods in deep RL have been developed and evaluated in mostly stationary settings, and it is often assumed that these solution methods are naturally applicable to continual learning. In this paper, we showed that canonical value-based deep RL methods like DQN and Rainbow are, in fact, not able to perform well in continual learning problems.

Our experiments were performed in a non-stationary setting where the agent is tasked to play a sequence of Atari 2600 games without any resetting in between. The results demonstrated that deep RL algorithms can lose their ability to learn (i.e. update the weights of the neural network), a phenomenon known as the loss of plasticity. We inspected the evolution of weights, gradients, and activations over the course of learning and observed activation collapse: in the continual setup with a non-stationary environment, a tiny fraction of network units produce non-zero activation value, inhibiting adaptation. Finally, we demonstrated that swapping ReLUs with CReLUs mitigate the loss of plasticity.

There are several promising directions for future work. Of course, CReLUs are not the only way of tackling the issues identified in this paper. Given the rationale for using CReLU activations, it would be interesting to verify whether other alternative activation functions such as Leaky-ReLU (Mass et al., 2013) or SeLU (Scaled Exponential Linear Unit) (Klambauer et al., 2017) would also be able to prevent the observed loss of plasticity. Promising alternatives beyond changing the activation function include shrink and perturb (Ash & Adams, 2020), introducing new random parameters to assist the existing ones (Nikishin et al., 2023), re-initializing parts of the network randomly or according to some notion of utility (e.g., Dohare et al., 2021; Sokar et al., 2023). However, the choice of activation function will have a limited impact on catastrophic interference. Effective reuse of past experience—learning what to forget and what to preserve—is not likely to be resolved by a new activation function. This paper represents a small step toward the goal continual deep reinforcement learning: developing agents that can maintain plasticity in the face of an ever changing world and efficiently leverage prior learning.

ACKNOWLEDGEMENTS

This work was developed while the authors were at DeepMind. The authors would like to thank Thomas Degris for his thorough feedback on an earlier draft, and Joshua Davidson, Finbarr Timbers, and Doina Precup for useful discussions.

REFERENCES

- Jordan T. Ash and Ryan P. Adams. On warm-starting neural network training. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- 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.
- Marc G. Bellemare, Will Dabney, and Rémi Munos. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2017.
- Marc G. Bellemare, Salvatore Candido, Pablo Samuel Castro, Jun Gong, Marlos C. Machado, Subhodeep Moitra, Sameera S. Ponda, and Ziyu Wang. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature*, 588(7836):77–82, 2020.
- Tudor Berariu, Wojciech Czarnecki, Soham De, Jörg Bornschein, Samuel L. Smith, Razvan Pascanu, and Claudia Clopath. A study on the plasticity of neural networks. *CoRR*, abs/2106.00042, 2021.
- Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan D. Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, Craig Donner, Leslie Fritz, Cristian Galperti, Andrea Huber, James Keeling, Maria Tsimpoukelli, Jackie Kay, Antoine Merle, Jean-Marc Moret, Seb Noury, Federico Pesamosca, David Pfau, Olivier Sauter, Cristian Sommariva, Stefano Coda, Basil Duval, Ambrogio Fasoli, Pushmeet Kohli, Koray Kavukcuoglu, Demis Hassabis, and Martin A. Riedmiller. Magnetic control of Tokamak plasmas through deep reinforcement learning. *Nature*, 602:414–419, 2022.
- Shibhansh Dohare, A. Rupam Mahmood, and Richard S. Sutton. Continual backprop: Stochastic gradient descent with persistent randomness. *CoRR*, abs/2108.06325, 2021.

- Jesse Farebrother, Marlos C. Machado, and Michael Bowling. Generalization and regularization in DQN. In *NeurIPS Deep Reinforcement Learning Workshop*, 2018.
- Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Mohammadamin Barekatin, Alexander Novikov, Francisco J. R. Ruiz, Julian Schrittwieser, Grzegorz Swirszcz, David Silver, Demis Hassabis, and Pushmeet Kohli. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature*, 610(7930):47–53, 2022.
- Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Matteo Hessel, Ian Osband, Alex Graves, Volodymyr Mnih, Rémi Munos, Demis Hassabis, Olivier Pietquin, Charles Blundell, and Shane Legg. Noisy networks for exploration. In *International Conference on Learning Representations (ICLR)*, 2018.
- Robert M French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135, 1999.
- Raia Hadsell, Dushyant Rao, Andrei A. Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in Cognitive Sciences*, 24(12):1028–1040, 2020.
- Matteo Hessel, Joseph Modayil, Hado van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Gheshlaghi Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. Towards continual reinforcement learning: A review and perspectives. *Journal of Artificial Intelligence Research*, 75:1401–1476, 2022.
- Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. In *Neural Information Processing Systems (NeurIPS)*, 2017.
- Aviral Kumar, Rishabh Agarwal, Dibya Ghosh, and Sergey Levine. Implicit under-parameterization inhibits data-efficient deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2021.
- Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz Rodríguez. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information Fusion*, 58:52–68, 2020.
- Vincent Liu, Adam White, Hengshuai Yao, and Martha White. Measuring and mitigating interference in reinforcement learning. *Openreview.net*, 2020.
- Jerry Luo, Cosmin Paduraru, Octavian Voicu, Yuri Chervonyi, Scott Munns, Jerry Li, Crystal Qian, Praneet Dutta, Jared Quincy Davis, Ningjia Wu, Xingwei Yang, Chu-Ming Chang, Ted Li, Rob Rose, Mingyan Fan, Hootan Nakhost, Tinglin Liu, Brian Kirkman, Frank Altamura, Lee Cline, Patrick Tonker, Joel Gouker, Dave Uden, Warren Buddy Bryan, Jason Law, Deeni Fatiha, Neil Satra, Juliet Rothenberg, Molly Carlin, Satish Tallapaka, Sims Witherspoon, David Parish, Peter Dolan, Chenyu Zhao, and Daniel J. Mankowitz. Controlling commercial cooling systems using reinforcement learning. *CoRR*, abs/2211.07357, 2022.
- Clare Lyle, Mark Rowland, and Will Dabney. Understanding and preventing capacity loss in reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 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.
- Dhruv Madeka, Kari Torkkola, Carson Eisenach, Dean P. Foster, and Anna Luo. Deep inventory management. *CoRR*, abs/2210.03137, 2022.
- Amol Mandhane, Anton Zhernov, Maribeth Rauh, Chenjie Gu, Miaosen Wang, Flora Xue, Wendy Shang, Derek Pang, Rene Claus, Ching-Han Chiang, Cheng Chen, Jingning Han, Angie Chen, Daniel J. Mankowitz, Jackson Broshear, Julian Schrittwieser, Thomas Hubert, Oriol Vinyals, and Timothy A. Mann. MuZero with self-competition for rate control in VP9 video compression. *CoRR*, abs/2202.06626, 2022.
- Andrew L. Mass, Awni Y. Hannun, and Andrew Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *International Conference on Machine Learning (ICML)*, 2013.

- Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Wenjie Jiang, Ebrahim M. Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Azade Nazi, Jiwoo Pak, Andy Tong, Kavya Srinivasa, William Hang, Emre Tuncer, Quoc V. Le, James Laudon, Richard Ho, Roger Carpenter, and Jeff Dean. A graph placement methodology for fast chip design. *Nature*, 594(7862):207–212, 2021.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Evgenii Nikishin, Max Schwarzer, Pierluca D’Oro, Pierre-Luc Bacon, and Aaron C. Courville. The primacy bias in deep reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2022.
- Evgenii Nikishin, Junhyuk Oh, Georg Ostrovski, Clare Lyle, Razvan Pascanu, Will Dabney, and Andre Barreto. Deep reinforcement learning with plasticity injection. In *Reincarnating RL Workshop at the ICLR Conference*, 2023.
- Emmanouil Antonios Platanios, Abulhair Saparov, and Tom M. Mitchell. Jelly bean world: A testbed for never-ending learning. In *International Conference on Learning Representations (ICLR)*, 2020.
- Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gómez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Giménez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, Tom Eccles, Jake Bruce, Ali Razavi, Ashley Edwards, Nicolas Heess, Yutian Chen, Raia Hadsell, Oriol Vinyals, Mahyar Bordbar, and Nando de Freitas. A generalist agent. *Transactions on Machine Learning Research*, 2022.
- Andrei Alex Rusu, Sebastian Flennerhag, Dushyant Rao, Razvan Pascanu, and Raia Hadsell. Probing transfer in deep reinforcement learning without task engineering. In *Conference on Lifelong Learning Agents, (CoLLAs)*, 2022.
- Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay. In *International Conference on Learning Representations (ICLR)*, 2016.
- Wenling Shang, Kihyuk Sohn, Diogo Almeida, and Honglak Lee. Understanding and improving convolutional neural networks via concatenated rectified linear units. In *International Conference on Machine Learning (ICML)*, 2016.
- Ghada Sokar, Rishabh Agarwal, Pablo Samuel Castro, and Utku Evci. The dormant neuron phenomenon in deep reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2023.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT press, 2018.
- Richard S. Sutton, Anna Koop, and David Silver. On the role of tracking in stationary environments. In *International Conference on Machine Learning (ICML)*, 2007.
- Matthew E. Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(7), 2009.
- Sebastian Thrun. A lifelong learning perspective for mobile robot control. In *Intelligent Robots and Systems*, 1995.
- Han Wang, Erfan Miah, Martha White, Marlos C. Machado, Zaheer Abbas, Raksha Kumaraswamy, Vincent Liu, and Adam White. Investigating the properties of neural network representations in reinforcement learning. *CoRR*, abs/2203.15955, 2022.
- Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. Dueling network architectures for deep reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2016.
- Maciej Wołczyk, Michał Zajac, Razvan Pascanu, Łukasz Kuciński, and Piotr Miłoś. Continual world: A robotic benchmark for continual reinforcement learning. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2020.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning (ICML)*, 2017.
- Xianyuan Zhan, Haoran Xu, Yue Zhang, Xiangyu Zhu, Honglei Yin, and Yu Zheng. DeepThermal: Combustion optimization for thermal power generating units using offline reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2022.

A APPENDIX: ADDITIONAL RESULTS

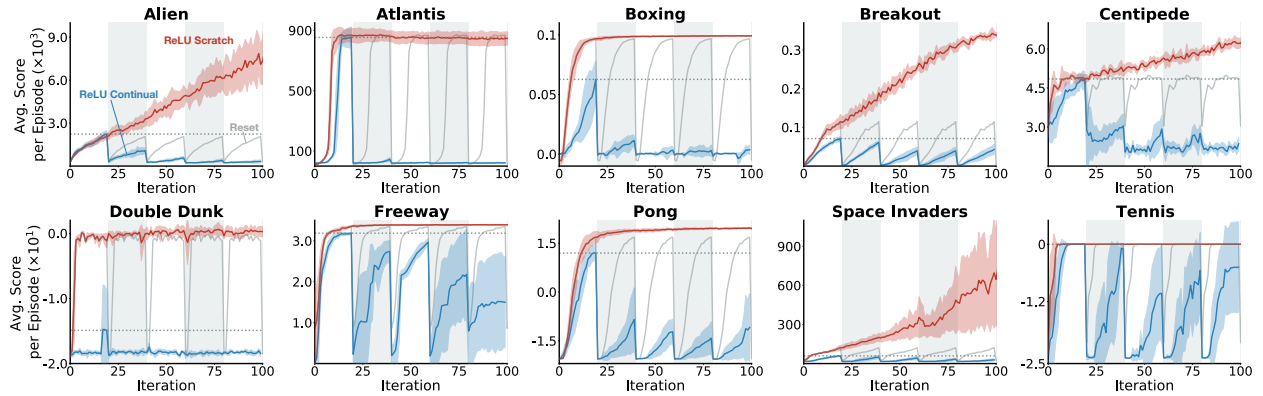


Figure 10: **Rainbow’s Loss of Plasticity on a repeating sequence of 10 Atari 2600 games with 20M steps on each game.** Each subplot reports learning within a single game. The blue curves show the Rainbow learning as it cycles through 10 games in a fixed sequence. The agent interacts with each game for 20M frames in every iteration of the cycle; the cycle restarts every 200M frames. The gray line depicts the performance achieved by an idealized reset (described in text) that represents the performance of an agent that is reset every 20M frames; there would be no impact from other games or previous visits to the current game. The dotted line marks the average performance achieved by the end of the first visit—a visual reference to highlight performance drop over successive visits. Rainbow’s performance when learning from scratch in a single game for 200M frames is represented in red. The red curve can be seen as an approximation to the idealized performance of an agent that does not suffer from interference (nor benefits from generalization) from other modes, and that does not forget anything about its past experiences, restarting where it left off in the previous visit.

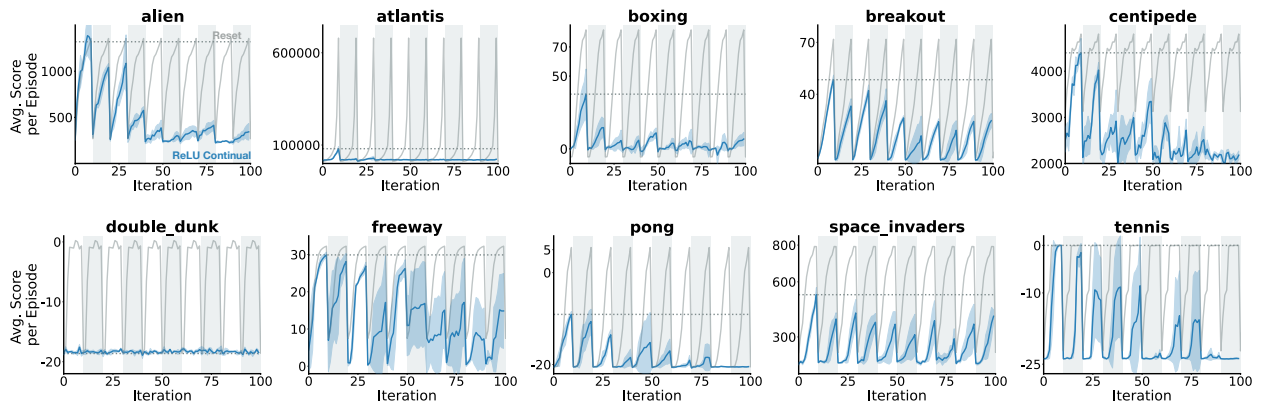


Figure 11: **Rainbow’s performance on a repeating sequence of 10 Atari games with 10M frames between game switches.** Each subplot reports learning within a single game. The blue curves show Rainbow learning as it cycles through 10 games in a fixed sequence. The agent interacts with each game for 10M frames in every iteration of the cycle; the cycle restarts every 100M frames. The gray line depicts the performance achieved by an idealized reset (described in text) that represents the performance of an agent that is reset every 10M frames; there would be no impact from other games or previous visits to the current game. The dotted line marks the average performance achieved by the end of the first visit—a visual reference to highlight performance drop over successive visits.

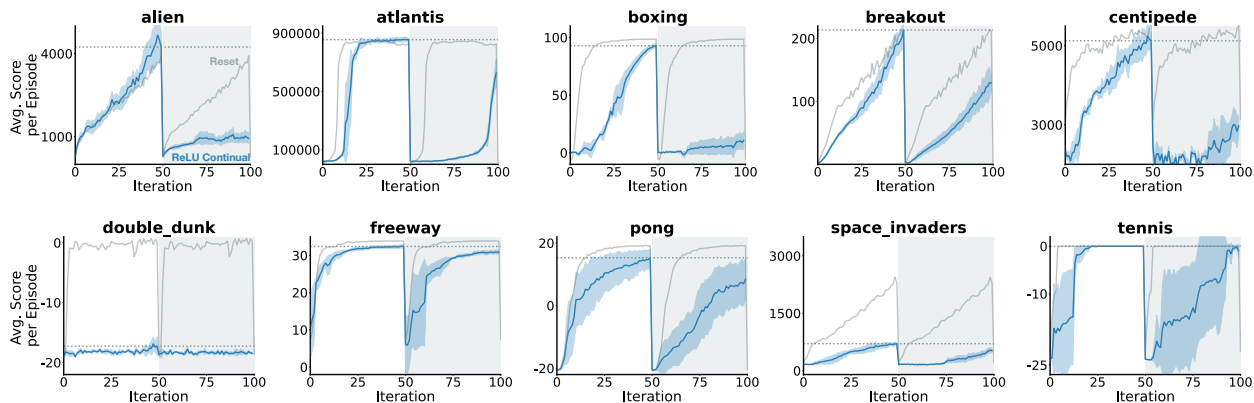


Figure 12: **Rainbow’s performance on a repeating sequence of 10 Atari games with 50M frames between game switches.** Each subplot reports learning within a single game. The blue curves show Rainbow learning as it cycles through 10 games in a fixed sequence. The agent interacts with each game for 50M frames in every iteration of the cycle; the cycle restarts every 500M frames. The gray line depicts the performance achieved by an idealized reset (described in text) that represents the performance of an agent that is reset every 50M frames; there would be no impact from other games or previous visits to the current game. The dotted line marks the average performance achieved by the end of the first visit—a visual reference to highlight performance drop over successive visits.

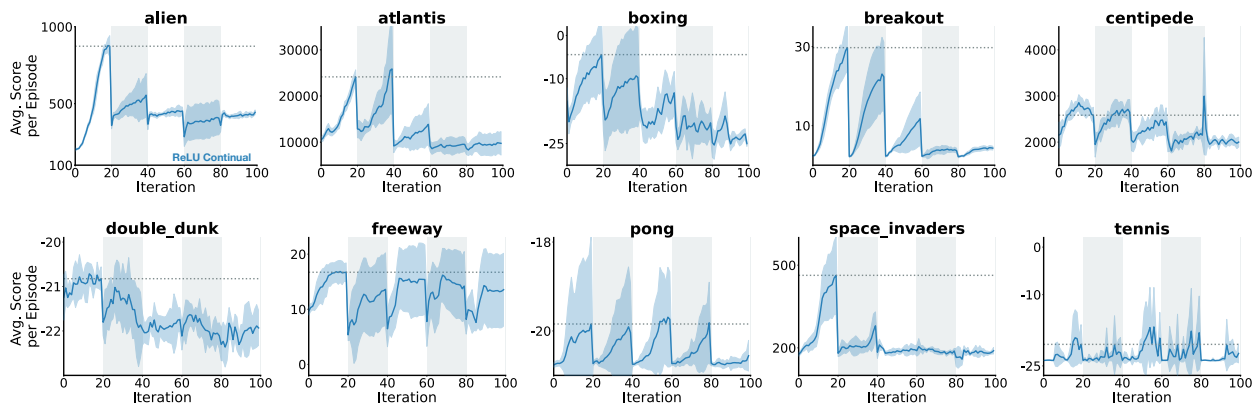


Figure 13: **DQN’s Loss of Plasticity on a repeating sequence of 10 Atari 2600 games with 20M steps on each game.** Each subplot reports learning within a single game. The blue curves show DQN learning as it cycles through 10 games in a fixed sequence. The agent interacts with each game for 20M frames in every iteration of the cycle; the cycle restarts every 200M frames. The dotted line marks the average performance achieved by the end of the first visit—a visual reference to highlight performance drop over successive visits.

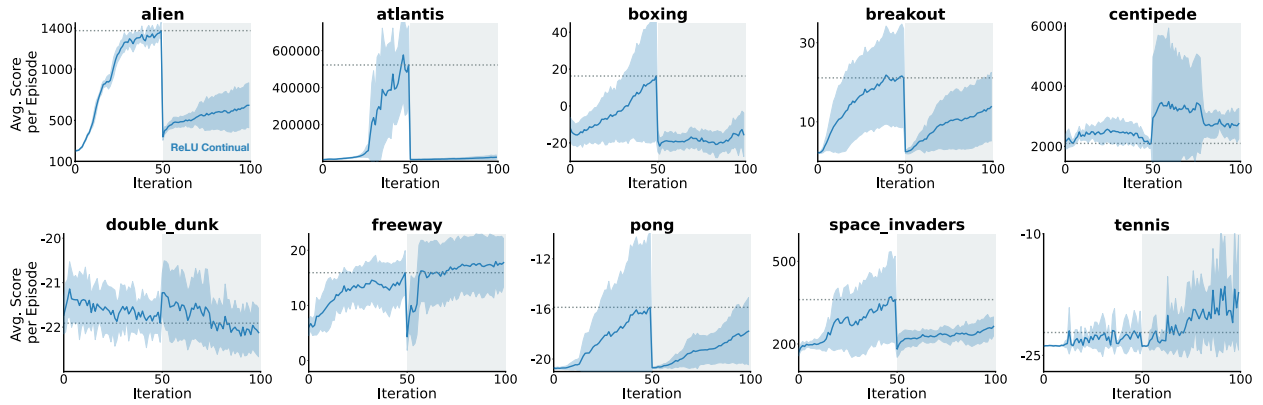


Figure 14: **DQN’s Loss of Plasticity on a repeating sequence of 10 Atari 2600 games with 50M steps on each game.** Each subplot reports learning within a single game. The blue curves show DQN learning as it cycles through 10 games in a fixed sequence. The agent interacts with each game for 50M frames in every iteration of the cycle; the cycle restarts every 500M frames. The dotted line marks the average performance achieved by the end of the first visit—a visual reference to highlight performance drop over successive visits.

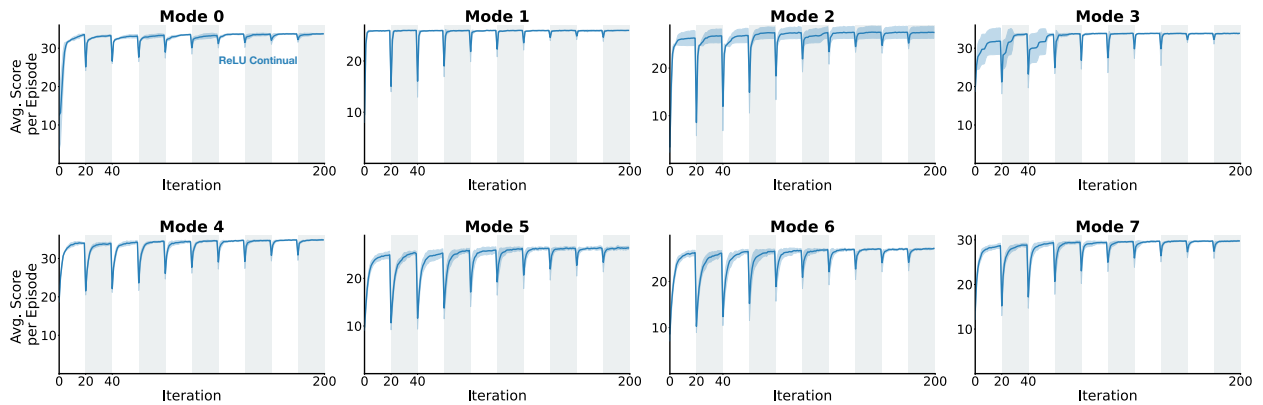


Figure 15: **Loss of Plasticity was not observed in Freeway on a sequence of 10 game modes.** The structure of this plot is identical to Figure 3. The blue curves show Rainbow learning as it cycles through 8 different game modes of Freeway in a fixed sequence. The agent interacts with each game mode for 20M frames in every iteration of the cycle; the cycle restarts every 160M frames.

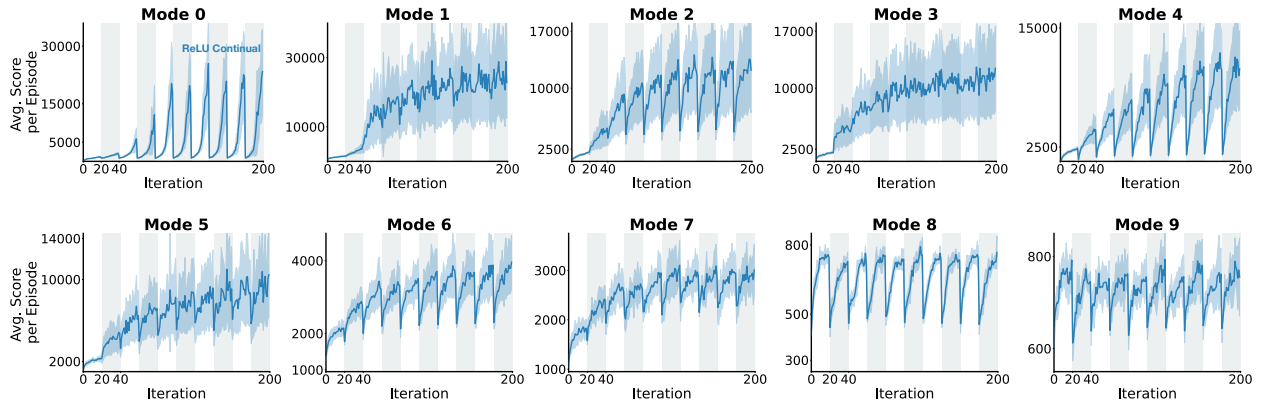


Figure 16: **Loss of Plasticity was not observed in Space Invaders on a sequence of 10 game modes.** The structure of this plot is identical to Figure 3. The blue curves show Rainbow learning as it cycles through 10 different game modes of Space Invaders in a fixed sequence. The agent interacts with each game mode for 20M frames in every iteration of the cycle; the cycle restarts every 200M frames.

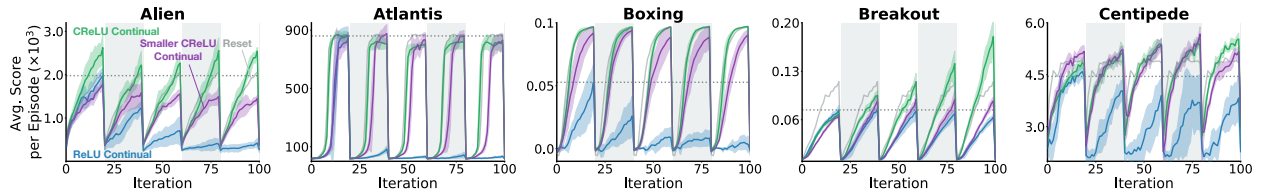


Figure 17: **Comparison of the performance induced by CReLUs with networks of different sizes.** As discussed in the main text, most plots in the paper were obtained with an invariant input dimension CReLU agent. For this CReLU variant the agent ends up with more parameters than the ReLU agent. Nevertheless, this difference is not the reason for CReLUs success. The purple curve depicts the performance of a Rainbow agent that also use CReLUs as activation functions, but the invariant output dimension one—consequently, this agent with CReLU activations has way fewer parameters than the Rainbow agent with ReLU activations, thus we termed it Smaler CReLU Continual. We can see that although the performance depicted in purple is worse than the one depicted in green, the purple curve, despite being generated by a neural network with way fewer parameters, still does not exhibit loss of plasticity.

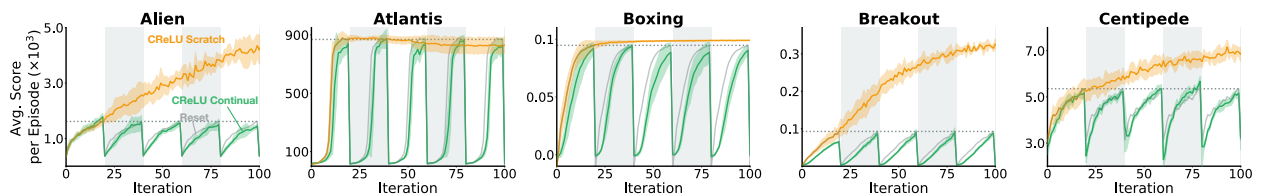


Figure 18: **CReLUs do not face loss of plasticity.** The orange curve depicts the learning progress of the Rainbow agent when learning from scratch in each game using CReLUs as activation functions, instead of ReLUs. We see the Rainbow agent cycling over multiple games consistently gets back to the final performance the Rainbow agent with CReLU activations reach at the end of the first iteration. Nevertheless, it is clear that the performance of the agent that cycles through multiple games is nowhere near the performance of an agent that learns from scratch in a single game. This plot also adds evidence that CReLUs performance is not due to a change in the neural network architecture as both agents in this plot have the same architecture.

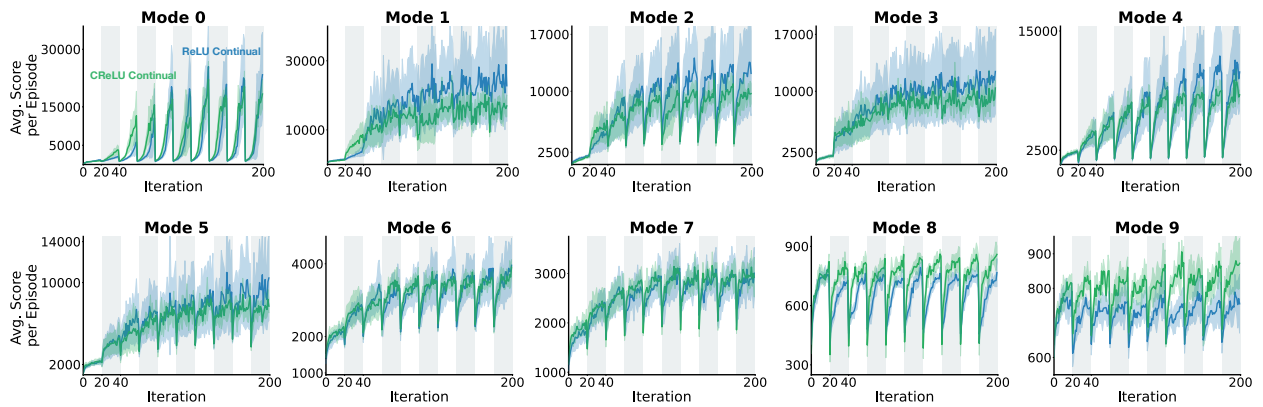


Figure 19: **Loss of plasticity was not observed in Space Invaders on a sequence of 10 game modes.** The structure of this plot is identical to Figure 16. The blue curves show the performance of the Rainbow agent with ReLU activations in the continual learning setting as it cycles through 10 different game modes of Space Invaders in a fixed sequence. The agent interacts with each game mode for 20M frames in every iteration of the cycle; the cycle restarts every 200M frames. The green curves show the performance of Rainbow with CReLU activations in the same continual learning setting. Neither agent exhibits loss of plasticity.

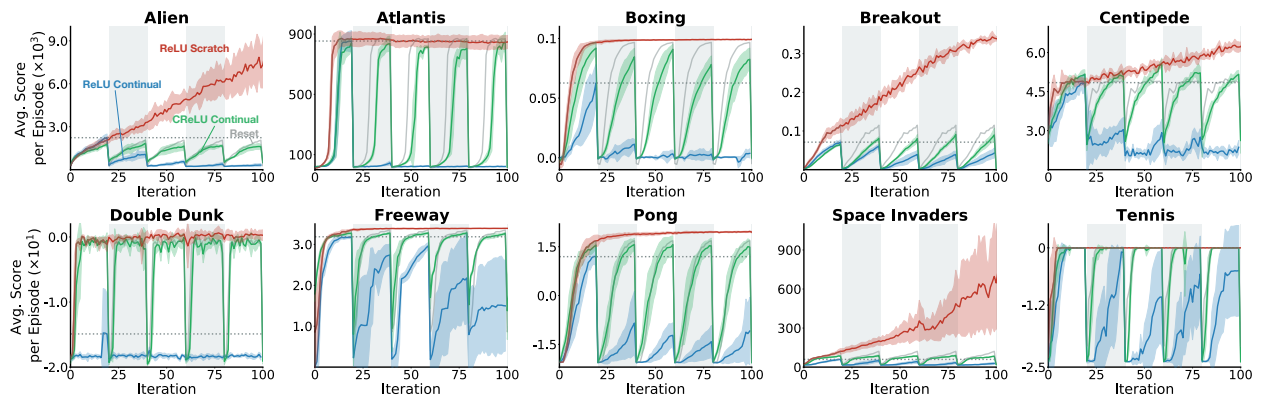


Figure 20: **Catastrophic forgetting alongside loss of plasticity and CReLU’s performance.** This figure reports the same data as Figure 10, with the addition of the learning progress of the Rainbow agent when using CReLUs as activation function (in green). The red curve can be seen as an approximation to the idealized performance of an agent that does not suffer from interference (nor benefits from generalization) from other modes and that does not forget anything about its past, restarting where it left off in the previous visit. Although an unrealistic and unfair baseline, it highlights how much Rainbow is leaving on the table when it cycles through several games, and how the performance achieved when using CReLUs, despite it often surpassing the level of performance the agent achieved the previous time it visited that game, it is still substantially worse than what one can achieve when learning only in a single game.