

THE EFFECTIVENESS OF WORLD MODELS FOR CONTINUAL REINFORCEMENT LEARNING

Samuel Kessler¹, Mateusz Ostaszewski², Michał Bortkiewicz², Mateusz Żarski³,
Maciej Wołczyk⁴, Jack Parker-Holder¹, Stephen J. Roberts¹, Piotr Miłoś⁵

¹ University of Oxford

² Warsaw University of Technology

³ Institute of Theoretical and Applied Informatics, PAS

⁴ Jagiellonian University, Cracow

⁵ Polish Academy of Sciences and Ideas NCBR

skessler@robots.ox.ac.uk

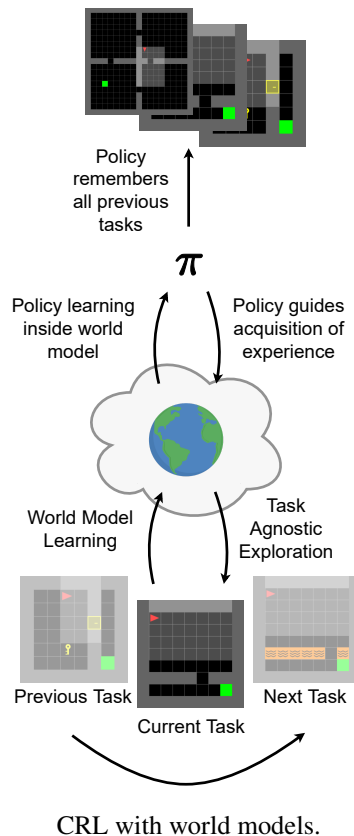
ABSTRACT

World models power some of the most efficient reinforcement learning algorithms. In this work, we showcase that they can be harnessed for continual learning – a situation when the agent faces changing environments. World models typically employ a replay buffer for training, which can be naturally extended to continual learning. We systematically study how different selective experience replay methods affect performance, forgetting, and transfer. We also provide recommendations regarding various modeling options for using world models. The best set of choices is called Continual-Dreamer, it is task-agnostic and utilizes the world model for continual exploration. Continual-Dreamer is sample efficient and outperforms state-of-the-art task-agnostic continual reinforcement learning methods on Minigrid and Minihack benchmarks.

1 INTRODUCTION

There have been many recent successes in reinforcement learning (RL), such as in games (Schrittwieser et al., 2019), robotics (OpenAI et al., 2018) and in scientific applications (Nguyen et al., 2021; Degraeve et al., 2022). However, these successes showcase methods for solving individual tasks. Looking beyond, the field of continual reinforcement learning (CRL) aims to develop agents which can solve many tasks, one after another, continually while retaining performance on all previously seen ones (Khetarpal et al., 2020). Such capabilities are conjectured to be essential for truly scalable intelligent systems, Ring (1994) and Hassabis et al. (2017) conjecture that truly scalable intelligent systems will additionally need to master many tasks in a continual manner. The field of continual reinforcement learning (CRL) aims to develop agents which can solve many tasks, one after another, continually while retaining performance on all previously seen ones (Khetarpal et al., 2020).

World models combine generative models with RL and have become one of the most successful paradigms in single-task RL (Ha & Schmidhuber, 2018). World models are typically trained iteratively, first, the policy interacts with the environment, second, the collected experience is used to train the world model and then the policy is trained using hallucinated experience generated by the world model (Kaiser et al., 2019). Such approaches are typically sample efficient, offloading the most data-hungry trial-and-error RL training to the imagined world. Further, generative models can be used to create compact representations that facilitate policy training and so obtain very good results for challenging pixel-based environments Hafner et al. (2023).



This paper explores using world models for learning tasks sequentially. We showcase a method that satisfies the traditional CL desiderata: avoids catastrophic forgetting, achieves transfer, high average performance, and is scalable. The proposed method *Continual-Dreamer* is built on top of DreamerV2 (Hafner et al., 2020) – an RL algorithm, which achieves state-of-the-art on a number of benchmarks. We define the Continual-Dreamer configuration in Section 5. DreamerV2 uses replay buffer (Lin, 1992), which we extend across many tasks to mitigate forgetting akin to (Isele & Cosgun, 2018; Rolnick et al., 2019). Additionally, we demonstrate that world models are capable to operate without explicit task identification. This is an important requirement in many CRL scenarios and enables further capabilities. In particular, we implement an adaptive exploration method similar to (Steinparz et al., 2022), which implicitly adapts to task changes. Based on our empirical results we argue that world-models offer a potent approach for CRL and should attract more research attention.

Our contributions are as follows¹:

- We present the first approach to task-agnostic model-based CRL. We use DreamerV2 as a backbone, and our work is transferable to other world models.
- We evaluate our method performance on two challenging CRL benchmarks, Minigrid and Minihack, and show that it outperforms state-of-the-art task-agnostic methods, demonstrating that the model-based paradigm is a viable solution to CRL.
- We thoroughly explore different experience replay strategies, which address how we sample from or populate the experience replay buffer to balance preventing forgetting of previously learned skills while enabling learning new skills from new tasks.

2 PRELIMINARIES

2.1 REINFORCEMENT LEARNING

A Partially Observable Markov Decision Process (POMDP (Kaelbling et al., 1998)) is the following tuple $(\mathcal{S}, \mathcal{A}, P, R, \Omega, \mathcal{O}, \gamma)$. Here, \mathcal{S} and \mathcal{A} are the sets of states and actions respectively, such that for $s_t, s_{t+1} \in \mathcal{S}$ and $a_t \in \mathcal{A}$. $P(s_{t+1}|s_t, a_t)$ is the transition distribution and $R(a_t, s_t, s_{t+1})$ is the reward function. Additionally, $\gamma \in (0, 1)$ is the discount factor. Since the environments we consider are partially observable, the agent does not have access to the environment state $s \in \mathcal{S}$, but only the observations $o \in \Omega$, where Ω is the set of observations and $\mathcal{O} : \mathcal{S} \times \mathcal{A} \rightarrow P(\Omega)$ is an observation function that defines a distribution over observations. Actions a_t are chosen using a policy π that maps observations to actions: $\Omega \rightarrow \mathcal{A}$. For the purposes of this introduction, let us assume we have access to the states s_t and we are working with a finite horizon H . Then the return from a state is $R_t = \sum_{i=t}^H \gamma^{(i-t)} r(s_i, a_i)$. In RL the objective is to maximize the expected return $J = \mathbb{E}_{a_i \sim \pi, s_0 \sim \rho} [R_1 | s_0]$ where $s_0 \sim \rho(s_0)$ and $\rho(\cdot)$ is the initial state distribution.

One approach to maximizing expected return is to use a *model-free* approach, to learn a policy $\pi_\phi : \mathcal{S} \rightarrow \mathcal{A}$ with a parametric model such as a neural network with parameters ϕ guided by an action-value function $Q_\theta(s_t, a_t)$ with parameters θ . Alternatively, instead of learning a policy directly from experience we can employ *model-based* RL (MBRL) and learn an intermediate model f , for instance, a transition model $s_{t+1} = f(s_t, a_t)$ from experience and learn our policy with additional experience generated from the model f (Sutton, 1991). Instead of working with the actual state s_t , our methods consider the observations o_t , we employ recurrent policies, action-value functions, and models to help better estimate states s_t (Hausknecht & Stone, 2015).

2.2 CONTINUAL REINFORCEMENT LEARNING

In continual RL the agent has a budget of N interactions with each task environment \mathcal{T}_τ . The agent is then required to learn a policy to maximize rewards in this environment, before interacting with a new environment and having to learn a new policy. Each task is defined as a new POMDP, $\mathcal{T}_\tau = (\mathcal{S}_\tau, \mathcal{A}_\tau, P_\tau, R_\tau, \Omega_\tau, \mathcal{O}_\tau, \gamma_\tau)$. See Appendix A.1 for a definition of CL in the supervised learning setting.

The agent is continually evaluated on all past and present tasks and so it is desirable for the agent’s policy to transfer to new tasks while not forgetting how to perform past tasks. CRL is not a new problem setting (Thrun & Mitchell, 1995), however, its definition has evolved over time and some settings will differ from paper to paper, we employ the setting above which is related to previous recent work in CRL (Kirkpatrick et al., 2016; Schwarz et al., 2018; Rolnick et al., 2019; Wolczyk et al., 2021; Kessler et al., 2021; Powers et al., 2021; Caccia et al., 2022).

¹Code available: <https://github.com/skezle/continual-dreamer>

Assumption on the task distribution. In this work we set out to study how world models deal with changing states spaces and transition distributions. We assume that for all pairs of tasks \mathcal{T}_i and \mathcal{T}_j that the state-spaces between tasks are disjoint: $\forall(\mathcal{T}_i, \mathcal{T}_j), \mathcal{S}_i \cap \mathcal{S}_j = \emptyset$.² This setting is popular for evaluating CRL benchmarks (Wolczyk et al., 2021; Powers et al., 2021).

3 RELATED WORK

Continual Reinforcement Learning. Seminal work in CRL, EWC (Kirkpatrick et al., 2016) enables DQN (Mnih et al., 2015) to continually learn to play different Atari games with limited forgetting. EWC learns new Q-functions by regularizing the parameter-weighted L2 distance between the new task’s current weights and the previous task’s optimal weights. EWC requires additional supervision informing it of task changes to update its objective, select specific Q-function head and select a task-specific ϵ -greedy exploration schedule. Progress and Compress (Schwarz et al., 2018) applies a regularization to policy and value function feature extractors for an actor-critic approach. Alternatively, LPG-FTW (Mendez et al., 2020) learns an actor-critic that factorizes into task-specific parameters and shared parameters. Both methods require task supervision and make use of task-specific parameters and shared parameters. Task-agnostic methods like CLEAR (Rolnick et al., 2019) do not require task information to perform CRL. CLEAR leverages experience replay buffers (Lin, 1992) to prevent forgetting: by using an actor-critic with V-trace importance sampling (Espeholt et al., 2018) of past experiences from the replay buffer. Model-based RL approaches to CRL have been demonstrated where the model weights are generated from a hypernetwork which itself is conditioned by a task embedding (Huang et al., 2021). Recent work demonstrates that recurrent policies for POMDPs can obtain good overall performance on continuous control CRL benchmarks (Caccia et al., 2022). Another task-aware solution is to expand a subspace of policies, the number of policies scaling sublinearly with the number of tasks (Gaya et al., 2022). Of the related works presented, only CLEAR is task-agnostic and so is the primary baseline under consideration when comparing task-agnostic world model approaches to CRL. A number of previous works have studied transfer in multi-task RL settings where the goals within an environment change (Barreto et al., 2017; Schaul et al., 2015; Barreto et al., 2019). In particular, incorporating the task definition directly into the value function (Schaul et al., 2015) and combining this with off-policy learning allows a CRL agent to solve multiple tasks continually, and generalize to new goals (Mankowitz et al., 2018). Model-based RL methods have also been assessed by looking at how quickly they can adapt to changes in reward function (Van Seijen et al., 2020). When only the reward function changes between tasks then experience from the replay buffer can interfere to prevent learning a new task or cause forgetting (Wan et al., 2022). The problem of interference can be mitigated by having separate policy heads per task (Kessler et al., 2021). See Appendix A.2 for related works on continual supervised learning.

Continual Adaptation. Instead of focusing on remembering how to perform all past tasks, another line of research investigates quick adaptation to changes in the environment. This can be captured by using a latent variable and off-policy RL (Xie et al., 2020). Alternatively, one can meta-learn a model such that it can then adapt quickly to new changes in the environment (Nagabandi et al., 2018). All these works use small environment changes such as modification of the reward function or variations in gravity or mass of certain agent limbs as new tasks. The tasks which we consider in this work contain substantially different \mathcal{A}, \mathcal{S} from one task to the next. For example, skills such as opening doors with keys or avoiding lava, or crossing a river which are quite different in comparison. Continual exploration strategies which use curiosity (Pathak et al., 2017) can be added as an intrinsic reward in the face of non-stationary environments in infinite horizon MDPs (Steinparz et al., 2022). The disagreement between an ensemble of models that predict the next state from the current state and action has been shown to be effective for exploration (Pathak et al., 2019). Our proposed model uses Plan2Explore which uses forward prediction disagreement and outperforms curiosity-based methods (Sekar et al., 2020). The tasks themselves can be meta-learned using a latent variable world model and task similarities can be exploited when learning a new task (Fu et al., 2022).

Curriculum Learning. Another related area of research is open-ended learning which aims to build agents that generalize to unseen environments through a curriculum that starts off with easy tasks and then progresses to harder tasks thereby creating agents which can generalize (Wang et al., 2019; Team et al., 2021; Parker-Holder et al., 2022).

4 WORLD MODELS FOR CONTINUAL REINFORCEMENT LEARNING

We leverage world models for learning tasks sequentially without forgetting. We use DreamerV2 (Hafner et al., 2020) which introduces a discrete stochastic and recurrent world model that is state of the art on numerous single-GPU RL benchmarks. This is a good choice for CRL since the world model is trained by reconstructing state, action, and reward

²Since we are working with POMDPs certain observations from different tasks might coincide.

trajectories from experience, we can thus leverage experience replay buffers which persist across tasks to prevent *forgetting* in the world model. Additionally, we can train a policy in the imagination or in the generated trajectories of the world model, similar to generative experience replay methods in supervised CL which remember previous tasks by replaying generated data (Shin et al., 2017). Thus, using a world model is also sample efficient. Also, world models are *task-agnostic* and do not require external supervision about task changes, without signaling to the agent that it is interacting with a new task. Additionally, by generating rollouts in the world model’s imagination the uncertainty in the world model’s predictions, more specifically the disagreement between predictions can be used as a task-agnostic exploration bonus. To summarize, we propose using model-based RL with recurrent world models as a viable method for CRL, see Algorithm 1 for an overview, with DreamerV2 as the world model. Recently, world models have been shown to collect diamonds in Minecraft, a very hard skill to achieve which requires the composition of many other skills with DreamerV3 (Hafner et al., 2023), the ideas introduced in this manuscript are directly applicable to newer world model methods as well.

Learning the World Model. DreamerV2 learns a recurrent (latent) state-space world model (RSSM) which predicts the forward dynamics of the environment. At each time step t the world model receives an observation o_t and is required to reconstruct the observations, o_t conditioned on the previous actions $a_{<t}$ (in addition to reconstructing rewards and discounts). The forward dynamics are modeled using an RNN, $h_t = \text{GRU}(h_{t-1}, z_t, a_t)$ (Chung et al., 2014) where h_t is the hidden state z_t are the discrete probabilistic latent states (Van Den Oord et al., 2017) which condition the observation predictions $p(o_t|z_t, h_t)$. Trajectories are sampled from an experience replay buffer and so persisting the replay buffer across different tasks should alleviate forgetting in the world model (Rolnick et al., 2019).

Algorithm 1 CRL with World Models

- 1: **Input:** Tasks (environments) $\mathcal{T}_{1:T}$, world model M , policy π , experience replay buffer \mathcal{D} .
 - 2: **for** \mathcal{T}_1 **to** \mathcal{T}_T **do**
 - 3: Train world model M on \mathcal{D} .
 - 4: Train π inside world model M .
 - 5: Execute π in task \mathcal{T}_τ to gather episodes and append to \mathcal{D} .
 - 6: **end for**
-

Policy Learning inside the World Model. The policy π is learned inside the world model by using an actor-critic (Sutton & Barto, 2018) while freezing the weights of the RSSM world model. At each step t of the dream inside the RSSM world model a latent state z_t is sampled, z_t , and the RNN hidden state condition the actor $\hat{a}_t \sim \pi(\cdot | z_t, h_t)$. The reward \hat{r}_{t+1} is predicted by the world model. The policy, π is then used to obtain new trajectories in the real environment. These trajectories are added to the experience replay buffer. An initial observation o_1 is used to start generating rollouts for policy learning. This training regime ensures that the policy generalizes to previously seen environments through the world model.

Task-agnostic Exploration. The policy learns using the imagined trajectories from the RSSM world model and the world model’s predicted rewards are used as a signal for the agent’s policy and critic. The policy is also used to gain experience inside the real environment. For exploration, the policy prioritizes regions of the state and action space where the world model produces uncertain predictions. Hence, the uncertainty in the world model’s trajectory prediction can be used as an additional intrinsic reward. This idea underpins Plan2Explore (Sekar et al., 2020) which naturally fits with DreamerV2.

The world model quantifies the uncertainty in the next latent state prediction by using a deep ensemble; multiple neural networks with independent weights. Deep ensembles are a surprisingly robust baseline for uncertainty quantification (Lakshminarayanan et al., 2017) and the ensemble’s variance is used as an intrinsic reward. The exploration neural networks in the ensemble are trained to predict the next RSSM latent features $[z_{t+1}, h_{t+1}]$. The world model is frozen while the ensemble is trained.

The policy π observes the reward $r = \alpha_i r_i + \alpha_e r_e$, where r_e is the extrinsic reward predicted by the world model, r_i is the intrinsic reward, the latent disagreement between the next latent state predictions. The coefficients α_i and α_e are $\in [0, 1]$. Hence the policy π can be trained inside the world model to seek regions in the state-action space that the world model struggles to predict and hence when the policy is deployed in the environment it will seek these same regions in the state-action space to obtain new trajectories to train the RSSM world model. The exploration strategy is significant for CRL since it is not task dependent unlike using DQN where each task needs an ϵ -greedy schedule (Kirkpatrick et al., 2016; Kessler et al., 2021) or SAC (Haarnoja et al., 2018) which needs an entropy regularizer per task (Wolczyk et al., 2021).

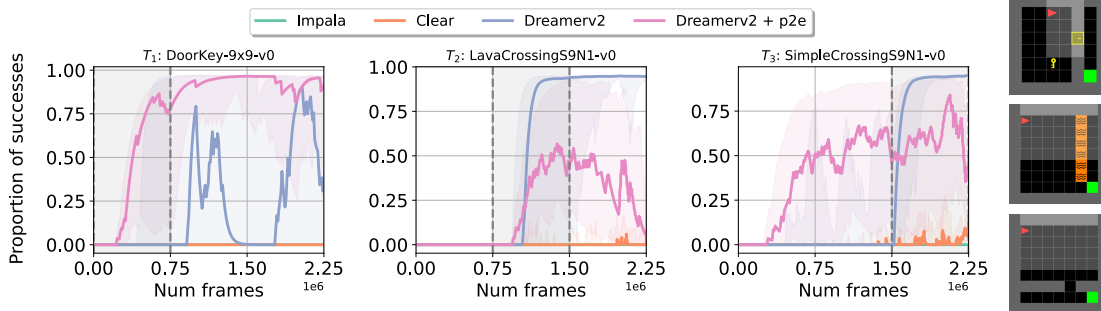


Figure 1: Performance of CRL agents on 3 Minigrid tasks. Grey-shaded regions indicate the environment which the agent is currently interacting with. All learning curves are a median and inter-quartile range across 20 seeds. On the right, we pick a random instantiation of the Minigrid environments that are being evaluated.

4.1 SELECTIVE EXPERIENCE REPLAY METHODS

To enable Continual-Dreamer to remember how to solve all the previous tasks it has learned with a limited replay buffer size requires us to select important trajectories to fill the experience replay buffer and selectively choose trajectories to train the world model. DreamerV2 uses a first-in-first-out (FIFO) replay buffer. It also randomly samples trajectories from the replay buffer to train the world model on and also randomly samples a trajectory so that the world model can start dreaming and allow the policy can learn inside the dream. In such a scenario, catastrophic forgetting can occur due to the loss of experience from previous tasks since the replay buffer is FIFO. There is prior work on managing experience replay buffers from supervised CL (Caccia et al., 2019) and off-policy RL (Isele & Cosgun, 2018) which we can systematically study for the application of CRL with world models. To ensure that we have a uniform sample of experience from all tasks in the replay buffer to sample from we explore the following methods:

- **Reservoir Sampling** (rs) (Vitter, 1985; Isele & Cosgun, 2018): enables a uniform distribution over all task experience seen so far. This is achieved by storing new examples in the replay buffer with a decreasing probability of $\min(n/t, 1)$, where t is the number of trajectories seen so far and n is the size of the replay buffer. This can be paired with any method of sampling from the replay buffer. By default, this is using random sampling. Reservoir sampling does not check for duplicate examples when storing experience. In practice, we found no duplicates when checking the episodes in the replay buffer using the experimental setup in Section 5.1, this is not a surprise since our task environments are procedurally generated, see Section 5 for further experimental details.
- **Coverage Maximization** (cm) (Isele & Cosgun, 2018): also attempts to create a uniform distribution of experience seen so far. Experience is added to the replay buffer by checking how close it is to trajectories already stored in the replay buffer. Trajectories are embedded using a fixed convolutional LSTM architecture (Shi et al., 2015) and we can calculate distances using an L^2 distance between the LSTM hidden state with respect to 1000 randomly selected trajectories from the replay buffer. The median distance determines the priority for the sample to be added to the replay buffer.

In addition to methods that populate the experience replay buffer, we can also consider how we should construct the mini-batch for world model and policy learning. For instance, we can prioritize more informative samples to help remembering and help learning new tasks to aid stability and plasticity. We consider 3 approaches:

- **Uncertainty sampling** (us): we construct a minibatch of experience where the probability of sampling an episode corresponds to the trajectory’s uncertainty or intrinsic reward from Plan2Explore. Next state uncertainties are generated for each transition and summed and normalized per trajectory before it is added to the replay buffer. We only calculate the uncertainty once before it is added to the replay buffer. This is similar to sampling the replay buffer according to the size of the temporal-difference error, known as sampling via “surprise” (Isele & Cosgun, 2018). The temporal difference error is also only calculated once when transitions are added to the experience replay buffer for DQN.
- **Reward sampling** (rwd) (Isele & Cosgun, 2018): we construct a mini-batch of experience for world model learning where the probability that an episode is sampled, corresponds to the reward from the environment.
- **50:50 sampling**, of past and recent experience. We construct a mini-batch for world model learning based on a 50:50 ratio of uniform random sampling from the replay buffer and sampling from a triangular distribution that favors the most recent experience added so far to help learning more recent tasks. This idea is similar to the on-policy off-policy ratio of learning in CLEAR (Rolnick et al., 2019) which also aims to balance stability and plasticity.

	Avg. Performance (\uparrow)	Avg. Forgetting (\downarrow)	Avg. Forward Transfer (\uparrow)
Impala	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
CLEAR	0.03 ± 0.05	0.01 ± 0.06	0.03 ± 0.03
Impala $\times 10$	0.16 ± 0.16	0.06 ± 0.13	-
CLEAR $\times 10$	0.64 ± 0.20	0.00 ± 0.00	-
DreamerV2	0.72 ± 0.24	-0.11 ± 0.30	0.49 ± 0.83
DreamerV2 + p2e	0.46 ± 0.10	0.05 ± 0.18	0.43 ± 0.22

Table 1: Results on 3 Minigrid tasks. All metrics are an average and standard deviation over 20 seeds. We use 0.75M interactions for each task and 7.5M in methods marked with $\times 10$. \uparrow indicates better performance with higher numbers, and \downarrow the opposite.

4.2 TASK-AWARE BASELINE

All replay buffer management techniques presented above are task-agnostic, i.e. operate without explicit task identification. We also consider a task-aware baseline for comparison, we use L^2 weight regularization with respect to the weights from the previous task, which is a simple regularization-based approach to CRL. In this scenario, after the first task, we add to each loss function, an additional objective that minimizes the distance between the current model and policy weights and the optimal weights from the previous task.

5 EXPERIMENTS

Our results indicate that DreamerV2 and DreamerV2 + Plan2Explore obtain good out-of-the-box performance for CRL on 3 Minigrid tasks (Chevalier-Boisvert et al., 2018). On a harder Minihack (Samvelyan et al., 2021) tasks from the CORA suite (Powers et al., 2021), we find that DreamerV2 and DreamerV2 + Plan2Explore exhibit forgetting. To address forgetting we systematically study various selective experience replay methods. The best configuration uses reservoir sampling (Vitter, 1985) which we name Continual-Dreamer and Continual-Dreamer + Plan2Explore. We use two primary baselines. First, Impala which is a powerful deep RL method not designed for CRL (Espeholt et al., 2018). Second, we consider CLEAR (Rolnick et al., 2019) which uses Impala as a base RL algorithm and leverages experience replay buffers to prevent forgetting and is task-agnostic.

Throughout our experiments, we use 3 different metrics average performance, average forgetting, and average forward transfer Appendix B, to assess the effectiveness of each method (Wolczyk et al., 2021) in addition to qualitatively inspecting the learning curves.

5.1 MINIGRID

We test the out-of-the-box performance of DreamerV2 and DreamerV2 + Plan2Explore as a CRL baseline on 3 sequential Minigrid tasks. Minigrid is a challenging image-based, partially observable, and sparse reward environment. The agent, in red, will get a reward of +1 when it gets to the green goal Figure 1. The agent sees a small region of the Minigrid environment as observation, o_t . We use 3 different tasks from Minigrid: DoorKey-9x9, SimpleCrossing-SN9 and LavaCrossing-9x9. Each environment has a different skill and so the tasks are diverse. Each method interacts with each task for 0.75M environment interactions, as previously proposed in (Kessler et al., 2021).

We evaluate CRL agents on all tasks, see Figure 1. The results indicate that DreamerV2 is able to solve difficult exploration tasks like the DoorKey-9x9. Additionally, since DreamerV2 trains its policy inside the world model it is more sample efficient than CLEAR which needs $\times 10$ more environment interactions to be able to solve the easier Minigrid tasks SimpleCrossing-SN9 and LavaCrossing-9x9, Table 1. The addition of Plan2Explore enables DreamerV2 to solve these environments even more quickly,

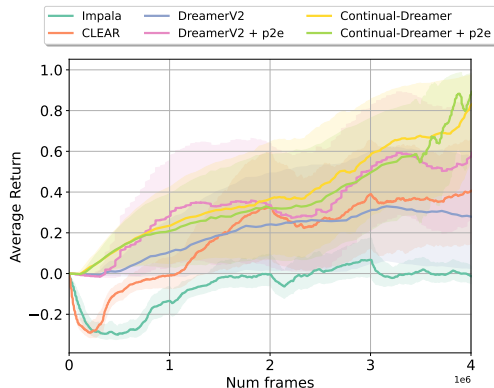


Figure 2: Return averaged over all tasks for various CRL agents on 4 Minihack tasks. All learning curves are IQM from the rliable package across 10 seeds and 1000 bootstrap samples (Agarwal et al., 2021).

see [Figure 1](#). DreamerV2 does exhibit some forgetting of the `DoorKey-9x9` task and this indicates that additional mechanisms to prevent forgetting might be needed.

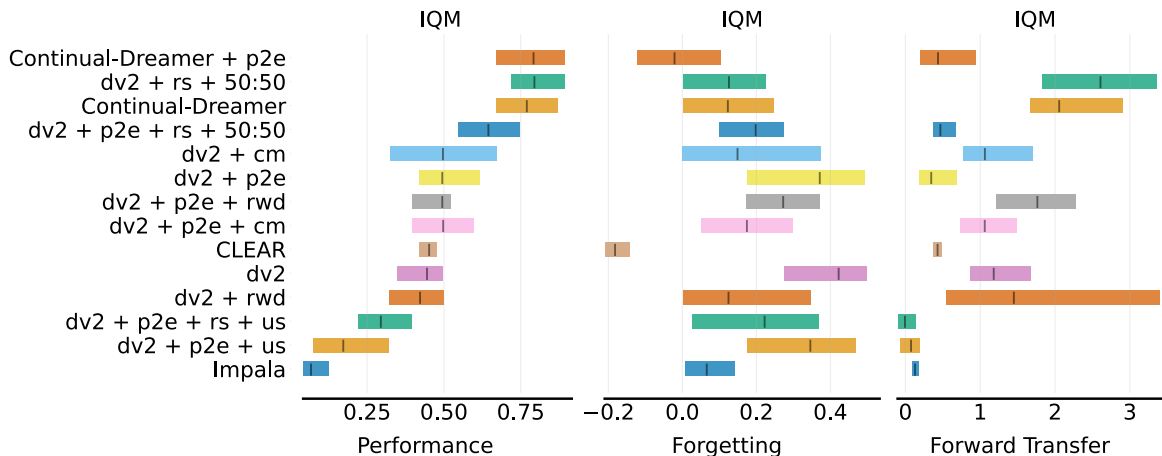


Figure 3: Metrics on 4 Minihack tasks using interquartile mean with 20 runs with different seeds and 1000 bootstrap samples from the `reliable` package ([Agarwal et al., 2021](#)).

From the metrics in [Table 1](#) we can see that DreamerV2 has strong forward transfer. From the learning curves for individual tasks [Figure 7](#) we can see that DreamerV2 struggles with independent task learning over the course of 1M environment steps. In contrast, when learning continually DreamerV2 is able to solve all tasks indicating that it transfers knowledge from previous tasks. This is not entirely surprising since the levels look similar and so the world model will be able to reconstruct observations of a new task more quickly compared to reconstruction from scratch.

For DreamerV2 we use the model and hyperparameters from ([Hafner et al., 2020](#)) with an experience replay buffer for world model learning of size 2M. For DreamerV2 + Plan2Explore we set the reward coefficients to $\alpha_i = \alpha_e = 0.9$ which was found by grid search of various single task Minihack environments over $\{0.1, 0.5, 0.9\}$ we use the same policy for exploration and evaluation and learn the world model by observation reconstruction only, rather than observation, reward, and discount reconstruction. We explore these design decisions using the Minihack benchmark in [Appendix D.4](#). For CLEAR we also use an experience replay buffer size of 2M like DreamerV2 and DreamerV2 + Plan2Explore.

5.2 MINIHACK

We test DreamerV2 and DreamerV2 with Plan2Explore on a set of harder Minihack tasks ([Samvelyan et al., 2021](#)). Minihack is a set of diverse, image-based, and sparse reward tasks based on the game of Nethack ([Küttler et al., 2020](#)) which have larger state spaces than MiniGrid and require learning more difficult skills such as crossing rivers by pushing multiple rocks into the river for instance. This will test the task-agnostic exploration mechanism from Plan2Explore further. We use 4 tasks Minihack tasks: in particular, we consider the following tasks `Room-Random-15x15-v0`, `Room-Trap-15x15-v0`, `River-Narrow-v0`, and `River-Monster-v0` which are a subset of the 12 Minihack tasks from the CORA CRL benchmark ([Powers et al., 2021](#)). Each task is seen once and has a budget of 1M environment interactions. We use the same experimental setup as in [Section 5.1](#), however, we keep the sizes of the experience replay buffer fixed to 1M for both DreamerV2, its variants, and CLEAR.

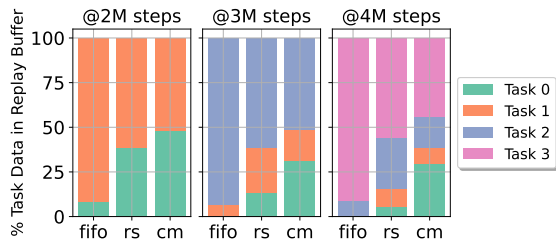


Figure 4: Proportion of episodes from each task in the replay buffer for different replay buffer construction strategies at 2M, 3M, and 4M environment steps on 4 task Minihack. Bar plots are an average of 5 runs.

We perform a comprehensive evaluation of various replay buffer management and mini-batch selection strategies. We find that using reservoir sampling together with DreamerV2 and DreamerV2 + Plan2Explore is the best configuration which we name Continual-Dreamer, see detailed analysis in [Section 5.2.1](#). The results of the Minihack experiment, as shown in [Figure 2](#), demonstrate that Continual-Dreamer and Continual-Dreamer + Plan2Explore perform better than the baselines regarding the average return over all tasks and 10 seeds. In particular, Continual-Dreamer and

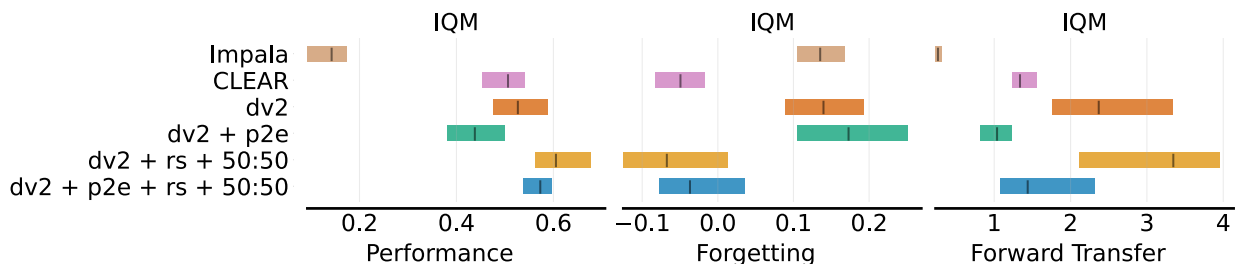


Figure 5: Metrics on 8 Minihack tasks using interquartile mean with 10 runs with different seeds and 1000 bootstrap samples. The results for DreamerV2 + rs + 50:50 with and without Plan2Explore are an interquartile mean with 5 runs with different seeds and 1000 bootstrap samples.

Continual-Dreamer + Plan2Explore exhibit faster learning on new tasks. Neither task-agnostic continual learning baselines, Impala and CLEAR, can learn the most difficult task *River-Monster-v0*, [Figure 9](#). These results suggest that world models are effective for consolidating knowledge across environments and changing task objectives. We also compare to a task-aware baseline: an L^2 regularization of the world model and actor-critic about the previous task’s optimal parameters [Figure 10](#). We find that this performs poorly.

5.2.1 EVALUATION OF DIFFERENT SELECTIVE REPLAY METHODS

[Figure 3](#) presents results for different replay strategies. If we consider DreamerV2 and DreamerV2 + Plan2Explore we can see that there is some forgetting of the first Room tasks, see [Figure 9](#). Our main finding is that reservoir sampling robustly mitigates forgetting, which gives rise to Continual-Dreamer.

We can increase plasticity by biasing the sampling toward recent experiences. We can see that if we add 50:50 sampling of the minibatch construction together with reservoir sampling, this causes inconsistent effects. For Continual-Dreamer, 50:50 sampling increases forgetting while the performance remains constant, indicating better learning of harder tasks and transfer. On the other hand, when 50:50 sampling is applied to Continual-Dreamer + Plan2Explore performance and forgetting worsen.

We tested DreamerV2’s performance in other variants, including uncertainty sampling (*us*), coverage maximization (*cm*), and reward sampling (*rdw*). The results we obtained are consistent with prior works ([Isele & Cosgun, 2018](#)). It can be seen that *us* performed closely to Impala with low performance, forward transfer, and high forgetting. Using *rdw* sampling results in performance that does not improve over random sampling of the minibatch. Using *cm* with DreamerV2 and DreamerV2 + Plan2Explore results in less forgetting. However, it behaves inconsistently when observing performance; performance increases when applied to DreamerV2 and it decreases when applied to DreamerV2 + Plan2Explore. As a baseline, we also tested a naive approach, which is to increase the size of the replay buffer; this indeed prevents forgetting and increases performance, see [Appendix D.5](#). However, this is at the cost of making it harder to learn new tasks: the harder exploration Minihack tasks are not solved and forward transfer decreases as we increase the size of the replay buffer. In [Figure 4](#) we can see the mini-batch task composition at 2M, 3M and 4M steps for 4 task Minihack for various replay buffer management methods. As training progresses the FIFO will lose experience from earlier tasks, whereas *rs* and *cm* are able to retain experience from earlier tasks. *rs* retains a more uniform distribution. Whereas *cm* - which is distance based - retains fewer samples from intermediate tasks, e.g. Task 1 at 4M steps since embeddings look similar to Task 0. We store entire episodes of state, action, and rewards in the replay buffer. When a task is mastered shorter episodes where the agent navigates to the goal are stored in the replay buffer. When adding experience from a new task the episodes are longer (episodes can reach the cut-off of 100 transitions before the episode is ended) and thus require more episodes of past data to be removed to make way for it in the reservoir sampling update. Due to this phenomenon, there are relatively fewer samples from the earlier tasks in the replay buffer. We perform a similar analysis for the sampling methods we consider in [Appendix D.8](#).

5.3 SCALING TO MORE TASKS

We scale to 8 Minihack tasks from the CORA suite: Room-Random-15x15-v0, Room-Monster-15x15-v0, Room-Trap-15x15-v0, Room-Ultimate-15x15-v0, River-Narrow-v0, River-v0, River-Monster-v0 and HideNSeek-v0. For DreamerV2 and its variants and for CLEAR we set the size of the experience replay buffer to 2M. We can see that by using reservoir sampling with 50:50 sampling prevents forgetting and slightly improves performance in DreamerV2 and DreamerV2 + Plan2Explore [Figure 5](#). Performance over CLEAR is also improved by introducing reservoir sampling and 50:50 sampling. The difficult

Minihack tasks, such as `River-v0`, `River-Monster-v0` and `HideNSeek-v0` are solved by DreamerV2 and its variants but not by CLEAR [Figure 11](#). We conjecture that warm-starting the world model by solving easier related tasks enables DreamerV2 to improve its downstream performance. For instance the single-task DreamerV2, see [Figure 8](#), does not solve `River-Monster-v0` as often as continual agent [Figure 11](#). Arguably, the skills learned in `Room-Monster-15x15-v0` and `River-v0` enable DreamerV2 to solve `River-Monster-v0`.

6 LIMITATIONS

Interference. Continual-Dreamer uses a replay buffer of past experience for retaining skills. Interference from past data is when past samples prevent the learning of new tasks in the face of changing reward functions or goal locations. This has been shown in world-models ([Wan et al., 2022](#)) and off-policy RL ([Kessler et al., 2021](#)). Continual-Dreamer and the selective experience replay methods which are explored are not intended to prevent interference from past samples. We show interference can affect DreamerV2 on the Minigrid `FourRooms-v0` environment with changing rewards/goals for 2 different tasks in [Appendix D.6](#).

Task data imbalances. The selective experience replay considered for managing experience in the replay buffer are not designed for keeping an even distribution of task data in the replay buffer in the face of imbalanced tasks. To illustrate this, we consider a 2 task scenario with 0.4M interactions of `Room-Random-15x15-v0` and then 2.4M interactions of `River-Narrow-v0` and a replay buffer of size 0.4M. We see that both *rs* and *cm* are unable to retain the performance of the first `Room-Random-15x15-v0` similarly to a FIFO buffer [Figure 6](#). For *rs* experience from the longer second task will saturate the replay buffer. While *cm* is able to retain a small number of samples from the first task, but not enough to prevent forgetting, since it uses a distanced-based criterion to populate the replay buffer. So *cm*-type approaches to replay buffer management could be a fruitful method for managing task imbalances.

7 DISCUSSION AND FUTURE WORKS

We have explored the use of world models as a *task-agnostic* CRL baseline. World models can be powerful CRL agents as they train the policy inside the world model and can thus be sample efficient. World models are trained by using experience replay buffers and so we can prevent *forgetting* of past tasks by persisting the replay buffer across from the current task to new tasks. Importantly, the world model’s prediction uncertainty can be used as an additional intrinsic task-agnostic reward to help exploration and solve difficult tasks in a task-agnostic fashion ([Sekar et al., 2020](#)). Previous CRL exploration strategies in the literature all require an indication of when the agent stops interacting with a previous task and starts interacting with a new task to reset an exploration schedule. Our implementation uses DreamerV2 as the world model ([Hafner et al., 2020](#)), and we demonstrate a selective experience replay setting which we call Continual-Dreamer which is a powerful CRL method on two difficult CRL benchmarks.

We show empirically that world models can be a strong task-agnostic baseline for CRL problems compared to state-of-the-art task-agnostic methods ([Rolnick et al., 2019](#)). DreamerV2 with Plan2Explore outperforms CLEAR on Minigrid. Our experiments on Minihack test the limits of using world models for CRL and require us to introduce experience replay buffer management methods to aid in retaining skills in addition to enabling the learning of new skills. We show that reservoir sampling enables an even coverage of experience in the replay buffer to mitigate forgetting. We call this configuration of DreamerV2 with reservoir sampling Continual-Dreamer. Future work will explore continuous control CRL benchmarks, such as ContinualWorld ([Wolczyk et al., 2021](#)) and explore other world-models.

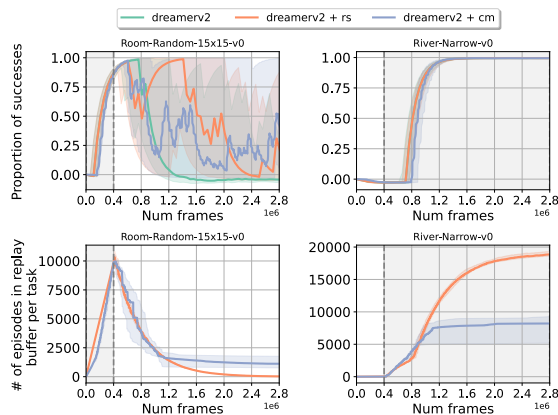


Figure 6: **Top**, learning curves for imbalanced number of interactions with `Room-Random-15x15-v0` and `River-Narrow-v0`. **Bottom**, the number of episodes from each task in the replay buffer for *rs* and *cm*. The grey shaded region indicates which timesteps the agent interacts with a task. All runs are interquartile ranges over 20 random seeds.

8 ACKNOWLEDGEMENTS

We would like to thank anonymous reviewers for their valuable feedback. SK would like to thank the Oxford-Man Institute of Quantitative Finance for funding. This research was funded by National Science Center Poland under the grant agreement 2020/39/B/ST6/01511 and from Warsaw University of Technology within the Excellence Initiative: Research University (IDUB) programme. We gratefully acknowledge Polish high-performance computing infrastructure PLGrid (HPC Centers: ACK Cyfronet AGH, CI TASK) for providing computer facilities and support within computational grant no. PLG/2023/016202. Piotr Milos was supported by the Polish National Science Centre grant 2019/35/O/ST6/03464.

REFERENCES

- Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in neural information processing systems*, 34:29304–29320, 2021.
- Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems*, 2019.
- André Barreto, Will Dabney, Rémi Munos, Jonathan J Hunt, Tom Schaul, Hado P van Hasselt, and David Silver. Successor features for transfer in reinforcement learning. *Advances in neural information processing systems*, 30, 2017.
- André Barreto, Diana Borsa, Shaobo Hou, Gheorghe Comanici, Eser Aygün, Philippe Hamel, Daniel K Toyama, Jonathan J Hunt, Shibl Mourad, David Silver, et al. The option keyboard: Combining skills in reinforcement learning. 2019.
- Ari S Benjamin, David Rolnick, and Konrad P Kording. Measuring and Regularizing Networks in Function Space. In *International Conference on Learning Representations*, 2019.
- Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020.
- Lucas Caccia, Rahaf Aljundi, Eugene Belilovsky, Massimo Caccia, Laurent Charlin, and Tinne Tuytelaars. Online continual learning with maximally interfered retrieval. *Advances in Neural Information Processing (NeurIPS)*, 2019.
- Massimo Caccia, Jonas Mueller, Taesup Kim, Laurent Charlin, and Rasool Fakoore. Task-agnostic continual reinforcement learning: In praise of a simple baseline. *arXiv preprint arXiv:2205.14495*, 2022.
- Arslan Chaudhry, Marcus Rohrbach Facebook, A I Research, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip H S Torr, and Marc Aurelio Ranzato. On Tiny Episodic Memories in Continual Learning. *arxiv.org:1902.10486*, 2019.
- Maxime Chevalier-Boisvert, Lucas Willems, and Suman Pal. Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>, 2018.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022.
- Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymir Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures. In *International Conference on Machine Learning*. 2018.
- Haotian Fu, Shangqun Yu, Michael Littman, and George Konidaris. Model-based lifelong reinforcement learning with bayesian exploration. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=6I3zJn9S1sb>.

- Jean-Baptiste Gaya, Thang Doan, Lucas Caccia, Laure Soulier, Ludovic Denoyer, and Roberta Raileanu. Building a subspace of policies for scalable continual learning. *arXiv preprint arXiv:2211.10445*, 2022.
- David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2018.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pp. 1861–1870. PMLR, 2018.
- Danijar Hafner, Timothy Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. *arXiv preprint arXiv:2010.02193*, 2020.
- Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.
- Demis Hassabis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. Neuroscience-inspired artificial intelligence. *Neuron*, 95(2):245 – 258, 2017. ISSN 0896-6273. doi: <https://doi.org/10.1016/j.neuron.2017.06.011>.
- Matthew Hausknecht and Peter Stone. Deep recurrent q-learning for partially observable mdps. In *2015 AAAI fall symposium series*, 2015.
- Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *arXiv preprint arXiv:1810.12488*, 2018.
- Yizhou Huang, Kevin Xie, Homanga Bharadhwaj, and Florian Shkurti. Continual model-based reinforcement learning with hypernetworks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 799–805. IEEE, 2021.
- David Isele and Akansel Cosgun. Selective experience replay for lifelong learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, et al. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019.
- Samuel Kessler, Jack Parker-Holder, Philip Ball, Stefan Zohren, and Stephen J Roberts. Same state, different task: Continual reinforcement learning without interference. *arXiv preprint arXiv:2106.02940*, 2021.
- Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. Towards continual reinforcement learning: A review and perspectives. *arXiv preprint arXiv:2012.13490*, 2020.
- James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *CoRR*, abs/1612.00796, 2016.
- Heinrich Küttler, Nantas Nardelli, Alexander H. Miller, Roberta Raileanu, Marco Selvatici, Edward Grefenstette, and Tim Rocktäschel. The NetHack Learning Environment. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- Soochan Lee, Junsoo Ha, Dongsu Zhang, and Gunhee Kim. A neural dirichlet process mixture model for task-free continual learning. *arXiv preprint arXiv:2001.00689*, 2020.
- Zhizhong Li and Derek Hoiem. Learning without Forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- Long-Ji Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Mach. Learn.*, 8(3–4):293–321, May 1992. ISSN 0885-6125. doi: 10.1007/BF00992699.

- David Lopez-Paz and Marc ' Aurelio Ranzato. Gradient Episodic Memory for Continual Learning. In *Advances in Neural Information Processing Systems*, 2017.
- Daniel J Mankowitz, Augustin Židek, André Barreto, Dan Horgan, Matteo Hessel, John Quan, Junhyuk Oh, Hado van Hasselt, David Silver, and Tom Schaul. Unicorn: Continual learning with a universal, off-policy agent. *arXiv preprint arXiv:1802.08294*, 2018.
- Jorge A Mendez, Boyu Wang, and Eric Eaton. Lifelong Policy Gradient Learning of Factored Policies for Faster Training Without Forgetting. In *Advances in Neural Information Processing Systems*, 2020.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharmashan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 2015. doi: 10.1038/nature14236.
- Anusha Nagabandi, Chelsea Finn, and Sergey Levine. Deep online learning via meta-learning: Continual adaptation for model-based rl. *arXiv preprint arXiv:1812.07671*, 2018.
- Cuong V. Nguyen, Yingzhen Li, Thang D. Bui, and Richard E. Turner. Variational continual learning. In *International Conference on Learning Representations*, 2018.
- V Nguyen, SB Orbell, Dominic T Lennon, Hyungil Moon, Florian Vigneau, Leon C Camenzind, Liuqi Yu, Dominik M Zumbühl, G Andrew D Briggs, Michael A Osborne, et al. Deep reinforcement learning for efficient measurement of quantum devices. *npj Quantum Information*, 7(1):1–9, 2021.
- OpenAI, Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Józefowicz, Bob McGrew, Jakub W. Pachocki, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, and Wojciech Zaremba. Learning dexterous in-hand manipulation. *CoRR*, abs/1808.00177, 2018.
- Jack Parker-Holder, Minqi Jiang, Michael Dennis, Mikayel Samvelyan, Jakob Foerster, Edward Grefenstette, and Tim Rocktäschel. Evolving curricula with regret-based environment design. *arXiv preprint arXiv:2203.01302*, 2022.
- Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, pp. 2778–2787. PMLR, 2017.
- Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. Self-supervised exploration via disagreement. In *International conference on machine learning*, pp. 5062–5071. PMLR, 2019.
- Sam Powers, Eliot Xing, Eric Kolve, Roozbeh Mottaghi, and Abhinav Gupta. Cora: Benchmarks, baselines, and metrics as a platform for continual reinforcement learning agents. *arXiv preprint arXiv:2110.10067*, 2021.
- Mark Bishop Ring. *Continual learning in reinforcement environments*. PhD thesis, University of Texas at Austin, 1994.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems 32*, pp. 350–360. 2019.
- Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *CoRR*, abs/1606.04671, 2016.
- Mikayel Samvelyan, Robert Kirk, Vitaly Kurin, Jack Parker-Holder, Minqi Jiang, Eric Hambro, Fabio Petroni, Heinrich Kuttler, Edward Grefenstette, and Tim Rocktäschel. Minihack the planet: A sandbox for open-ended reinforcement learning research. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. URL <https://openreview.net/forum?id=skFwlyefkWJ>.
- Tom Schaul, Daniel Horgan, Karol Gregor, and David Silver. Universal value function approximators. In *International conference on machine learning*, pp. 1312–1320. PMLR, 2015.
- Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy P. Lillicrap, and David Silver. Mastering atari, go, chess and shogi by planning with a learned model. *CoRR*, abs/1911.08265, 2019.

- Jonathan Schwarz, Wojciech Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4535–4544. PMLR, 2018.
- Ramanan Sekar, Oleh Rybkin, Kostas Daniilidis, Pieter Abbeel, Danijar Hafner, and Deepak Pathak. Planning to explore via self-supervised world models. In *International Conference on Machine Learning*, pp. 8583–8592. PMLR, 2020.
- Xingjian Shi, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems*, 28, 2015.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual Learning with Deep Generative Replay. In *Advances in Neural Information Processing Systems*, 2017.
- Christian Steinparz, Thomas Schmied, Fabian Paischer, Marius-Constantin Dinu, Vihang Patil, Angela Bitto-Nemling, Hamid Eghbal-zadeh, and Sepp Hochreiter. Reactive exploration to cope with non-stationarity in lifelong reinforcement learning. *arXiv preprint arXiv:2207.05742*, 2022.
- Richard S Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4): 160–163, 1991.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Open Ended Learning Team, Adam Stooke, Anuj Mahajan, Catarina Barros, Charlie Deck, Jakob Bauer, Jakub Sygnowski, Maja Trebacz, Max Jaderberg, Michael Mathieu, et al. Open-ended learning leads to generally capable agents. *arXiv preprint arXiv:2107.12808*, 2021.
- Sebastian Thrun and Tom M. Mitchell. Lifelong robot learning. *Robotics and Autonomous Systems*, 15(1): 25–46, 1995. ISSN 0921-8890. doi: [https://doi.org/10.1016/0921-8890\(95\)00004-Y](https://doi.org/10.1016/0921-8890(95)00004-Y). URL <https://www.sciencedirect.com/science/article/pii/092188909500004Y>. The Biology and Technology of Intelligent Autonomous Agents.
- Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019.
- Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- Harm Van Seijen, Hadi Nekoei, Evan Racah, and Sarath Chandar. The loca regret: a consistent metric to evaluate model-based behavior in reinforcement learning. *Advances in Neural Information Processing Systems*, 33:6562–6572, 2020.
- Jeffrey Scott Vitter. Random sampling with a reservoir. *ACM Trans. Math. Softw.*, 11:37–57, 1985.
- Yi Wan, Ali Rahimi-Kalahroudi, Janarthanan Rajendran, Ida Momennejad, Sarath Chandar, and Harm H Van Seijen. Towards evaluating adaptivity of model-based reinforcement learning methods. In *International Conference on Machine Learning*, pp. 22536–22561. PMLR, 2022.
- Rui Wang, Joel Lehman, Jeff Clune, and Kenneth O Stanley. Paired open-ended trailblazer (poet): Endlessly generating increasingly complex and diverse learning environments and their solutions. *arXiv preprint arXiv:1901.01753*, 2019.
- Maciej Wolczyk, Michal Zajac, Razvan Pascanu, Lukasz Kucinski, and Piotr Milos. Continual world: A robotic benchmark for continual reinforcement learning. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 28496–28510, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/ef8446f35513a8d6aa2308357a268a7e-Abstract.html>.
- Annie Xie, James Harrison, and Chelsea Finn. Deep reinforcement learning amidst lifelong non-stationarity. *arXiv preprint arXiv:2006.10701*, 2020.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual Learning Through Synaptic Intelligence. In *International Conference on Machine Learning*, 2017.

APPENDIX

TABLE OF CONTENTS

A	Continual Supervised Learning	14
A.1	Definition	14
A.2	Related Works	14
B	Continual Reinforcement Learning Metrics	14
B.1	Average Performance	14
B.2	Forgetting	14
B.3	Forward Transfer	15
C	Single Task experiments	15
D	Further Experiments	15
D.1	4 task Minihack	15
D.2	Task-aware world-model baseline	16
D.3	Scaling to 8 tasks	16
D.4	DreamerV2 Ablation Experiments	16
D.5	Stability versus Plasticity: Increasing the Size of the Replay Buffer	18
D.6	Interference	18
D.7	Decreasing the Replay Buffer Size	18
D.8	Analysis of Replay Buffer Sampling Methods	19

APPENDIX A CONTINUAL SUPERVISED LEARNING

A.1 DEFINITION

Continual Learning (CL) is a setting whereby a model must master a set of tasks sequentially while maintaining performance across all previously learned tasks. Other important objectives are to develop scalable CL methods which are able to transfer knowledge from previous tasks to aid the learning of new tasks, known as forward transfer. Traditionally, in supervised CL, the model is sequentially shown T tasks, denoted \mathcal{T}_τ for $\tau = 1, \dots, T$. Each task, \mathcal{T}_τ , is comprised of a dataset $\mathcal{D}_\tau = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_\tau}$ which a neural network is required to learn from. More generally, a task is denoted by a tuple comprised of the conditional and marginal distributions, $\{p_\tau(y|\mathbf{x}), p_\tau(\mathbf{x})\}$. After task τ , the model will lose access to the training dataset for \mathcal{T}_τ , however, its performance will be continually evaluated on all tasks \mathcal{T}_i for $i \leq \tau$. For a comprehensive review of CL scenarios see (Hsu et al., 2018; Van de Ven & Tolias, 2019).

A.2 RELATED WORKS

Here, we briefly describe CL methods. One approach to CL referred to as *regularization approaches* regularizes a NN’s weights to ensure that optimizing for a new task finds a solution that is “close” to the previous task’s (Kirkpatrick et al., 2016; Nguyen et al., 2018; Zenke et al., 2017). Working with functions can be easier than with NN weights and so task functions can be regularized to ensure that learning new function mappings are “close” across tasks (Li & Hoiem, 2017; Benjamin et al., 2019; Buzzega et al., 2020). By contrast, *expansion approaches* add new NN components to enable learning new tasks while preserving components for specific tasks (Rusu et al., 2016; Lee et al., 2020). *Memory approaches* replay data from previous tasks when learning the current task. This can be performed with a generative model (Shin et al., 2017). Or samples from previous tasks (*memories*) (Lopez-Paz & Ranzato, 2017; Aljundi et al., 2019; Chaudhry et al., 2019).

APPENDIX B CONTINUAL REINFORCEMENT LEARNING METRICS

We describe in detail how to calculate the continual reinforcement learning metrics used extensively throughout this manuscript.

B.1 AVERAGE PERFORMANCE

This measures how well a CRL method performs on all tasks at the end of the task sequence. The task performance is $p_\tau(t) = [-1, 1]$ for all $\tau < T$. Since we have a reward of +1 for completing the task and -1 for being killed by a monster or falling into lava. If each task is seen for N environment steps and we have T tasks and the τ -th task is seen over the interval of steps $[(\tau - 1) \times N, \tau \times N]$. The average final performance metric for our continual learning agent is defined as:

$$p(t_f) = \frac{1}{T} \sum_{\tau=1}^T p_\tau(t_f), \quad (\text{B.1})$$

where $t_f = N \times T$ is the final timestep.

B.2 FORGETTING

The average forgetting is the performance difference after interacting with a task versus the performance at the end of the final task. The average forgetting across all tasks is defined as:

$$F = \frac{1}{T} \sum_{\tau=1}^T F_\tau \quad \text{where} \quad F_\tau = p_\tau(\tau \times N) - p_\tau(t_f). \quad (\text{B.2})$$

The forgetting of the final T -th task is $F_T = 0$. If a CRL agent has better performance at the end of the task sequence compared to after τ -th task at time-step $\tau \times N$ then $F_\tau < 0$. Note that both the average performance and the forgetting metrics are functions of $p_\tau(t_f)$ so we expect anti-correlation between these two metrics.

B.3 FORWARD TRANSFER

The forward transfer is the difference in task performance during continual learning compared to the single task performance. The forward transfer is defined:

$$FT = \frac{1}{T} \sum_{\tau=1}^T FT_\tau \quad \text{where} \quad FT_\tau = \frac{\text{AUC}_\tau - \text{AUC}_{\text{ref}_\tau}}{1 - \text{AUC}_\tau}, \quad (\text{B.3})$$

where AUC denotes the area under the curve and is defined as:

$$\text{AUC}_\tau = \frac{1}{N} \int_{(\tau-1) \times N}^{\tau \times N} p_\tau(t) dt \quad \text{and} \quad \text{AUC}_{\text{ref}_\tau} = \frac{1}{N} \int_0^N p_{\text{ref}_\tau}(t) dt. \quad (\text{B.4})$$

$FT_\tau > 0$ means that the CRL agent achieves better performance on task τ during continual learning versus in isolation. So this metric measures how well a CRL agent transfers knowledge from previous tasks when learning a new task.

APPENDIX C SINGLE TASK EXPERIMENTS

To assess the forward transfer of DreamerV2 for CRL we need the performance of each task as a reference [Equation \(B.3\)](#). Single task learning curves for Minigrid are shown in [Figure 7](#) and single task learning curves for all Minihack tasks are shown in [Figure 8](#).

APPENDIX D FURTHER EXPERIMENTS

We introduce further experiments which are referenced in the main paper. In [Appendix D.1](#) we show learning curves for each individual task for 4 task Minihack. In [Appendix D.2](#) we show the results of the regularization based task-aware world model baseline. In [Appendix D.3](#) we show learning curves for each individual task from 8 task Minihack experimental setup. In [Appendix D.4](#) we explore various design choices required

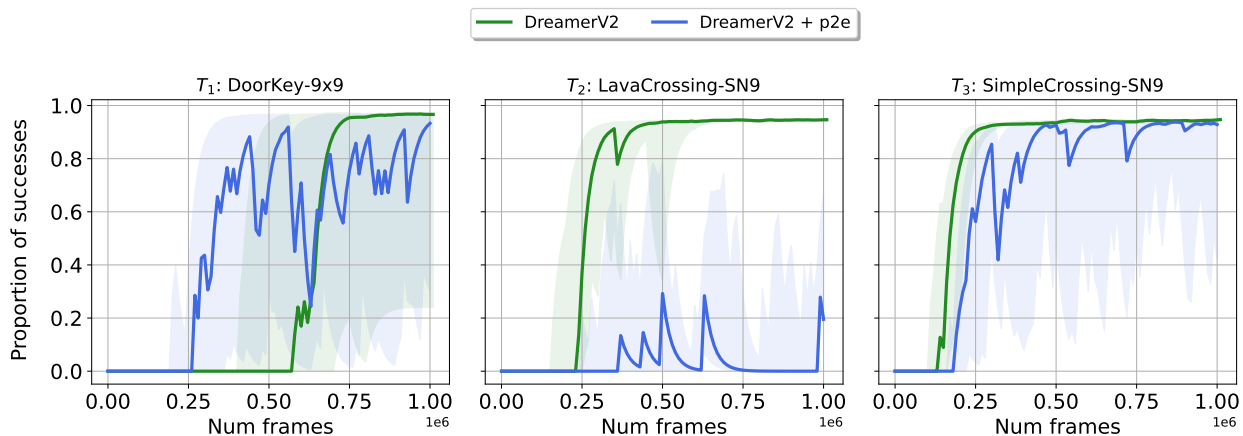


Figure 7: Single task performance of on individual tasks from the Minigrad CRL benchmark. All curves are a median and inter-quartile range over 20 seeds.




Figure 8: Single task performance on individual tasks from the Minihack CRL benchmark. All curves are a median and inter-quartile range over 10 seeds.

for DreamerV2 with Plan2Explore to get the best performance for CRL. In [Appendix D.5](#) we explore how increasing the size of the experience replay buffer size affects performance in the Minihack CRL benchmark. These experiments are on 8 tasks of Minihack. The 8 tasks are a subset of those introduced in the CORA Minihack suite ([Powers et al., 2021](#)) and are a superset of the 4 tasks in the main paper. The 8 tasks, in order, are: Room-Random-15x15-v0, Room-Monster-15x15-v0, Room-Trap-15x15-v0, Room-Ultimate-15x15-v0, River-Narrow-v0, River-v0, River-Monster-v0, and HideNSeek-v0.

D.1 4 TASK MINIHACK

[Figure 9](#) shows the success rates of the baseline methods and DreamerV2 variants for each of the four tasks in the Minihack experiment [Section 5.2](#). Continual Dreamer successfully balances retaining knowledge from previous tasks and learning new ones. On the other hand, the CLEAR baseline can only learn the first two tasks with a significant delay, while Impala struggles on all tasks. DreamerV2 and DreamerV2 + Plan2Explore perform poorly on the last task and exhibit more forgetting than Continual-Dreamer.

D.2 TASK-AWARE WORLD-MODEL BASELINE

Differences for task-agnostic and task-aware variants of DreamerV2 are shown in [Figure 10](#). The task-aware variant based on L^2 regularization, underperforms in comparison to task-agnostic methods. The plausible explanation is that L^2 regularization makes the method too rigid to efficiently learn the two last tasks because of the excessive influence of the first tasks. We optimize the size of the L^2 scaling from the set $\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10, 100\}$.

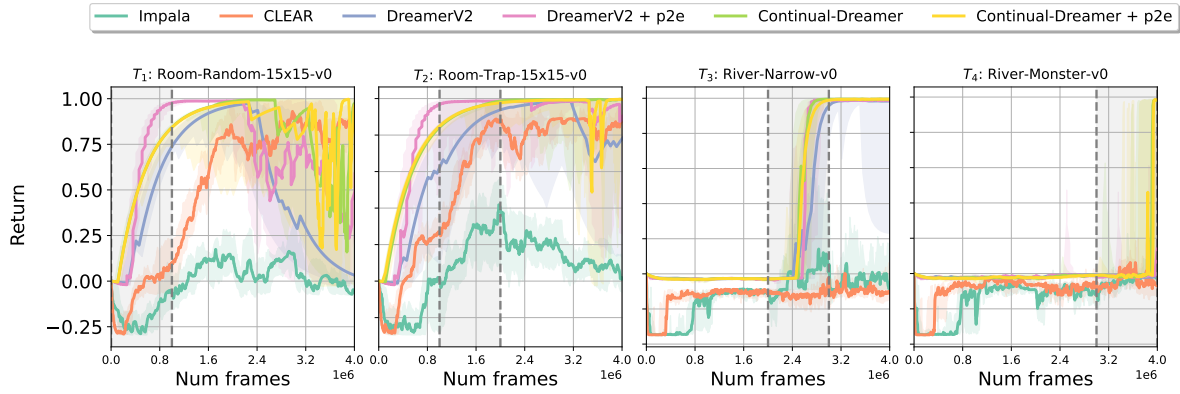


Figure 9: Detailed, per task, comparison of baselines and DreamerV2 variants that are presented in Figure 2 with average return. All curves are a median and inter-quartile range over 20 seeds.

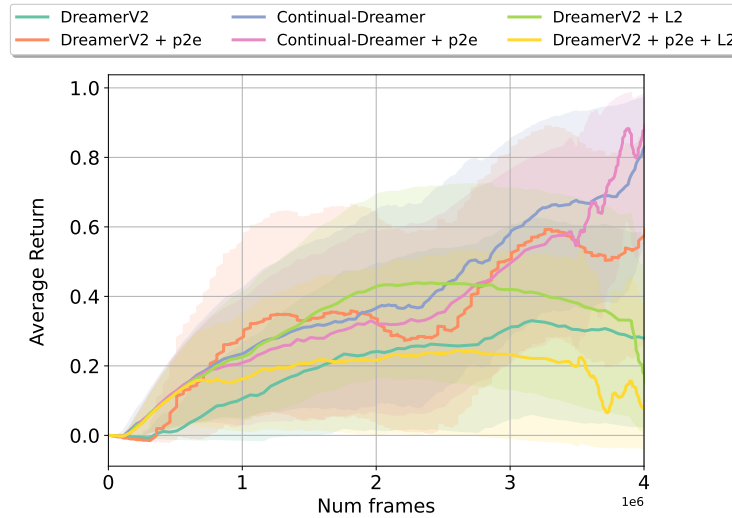


Figure 10: Comparison of the average return over all tasks between task-aware and task-agnostic approaches based on DreamerV2, on 4 Minihack tasks. All curves are IQM from `reliable` package across 10 seeds and 1000 bootstrap samples.




Figure 11: Learning curves for 8 Minihack tasks for DreamerV2 and variants and Impala and CLEAR baselines. All curves are a median and interquartile range of 10 seeds.

Plan2Explore	\hat{o} reconstruction only	$\pi_{exp} = \pi_{eval}$	Avg. Performance (\uparrow)	Avg. Forgetting (\downarrow)	Avg. Forward Transfer (\uparrow)
-	-	-	0.09 ± 0.07	0.37 ± 0.07	0.56 ± 0.86
✓	-	-	0.28 ± 0.13	0.13 ± 0.08	0.11 ± 0.15
✓	✓	-	0.39 ± 0.13	0.19 ± 0.16	0.87 ± 0.95
✓	✓	✓	0.38 ± 0.03	0.22 ± 0.05	0.76 ± 0.25

Table 2: CRL metrics for different design decisions on DreamerV2 for the Minihack CRL benchmark of 8 tasks. All metrics are an average and standard deviation over 5 seeds. \uparrow indicates better performance with higher numbers, and \downarrow the opposite.

D.3 SCALING TO 8 TASKS

The results for 8 Minihack tasks are shown in [Figure 11](#). DreamerV2 variants display strong knowledge retention and effective learning on almost every task compared to baseline methods. CLEAR method struggle with the last 4 tasks, whereas Impala’s performance is poor on every task. DreamerV2 and variants displays forgetting of the initial tasks, for which CLEAR retains the highest performance. However, CLEAR, in contrast to DreamerV2 variants, struggles to learn novel tasks.

D.4 DREAMERV2 ABLATION EXPERIMENTS

We explore various design choices which come from the implementations of DreamerV2 ([Hafner et al., 2020](#)) and Plan2Explore ([Sekar et al., 2020](#)).

1. The use of Plan2Explore as an intrinsic reward.
2. World model learning by reconstructing the observations \hat{o}_t only and not the observations, rewards, and discounts altogether.
3. The use of the exploration policy to evaluate the performance on all current and past tasks rather than having a separate exploration and evaluation policy.

The results are shown in [Table 2](#). We decided to pick the model in the final line in [Table 2](#) to report the results in the main paper as they produce good results on Minihack with a relatively small standard deviation.




Figure 12: CRL metric package for DreamerV2 + Plan2Explore for the Minihack benchmark of 8 tasks versus the experience replay buffer size of the world model for DreamerV2 + Plan2Explore. All metrics are an interquartile mean (IQM) over 5 seeds with 1000 bootstrap samples from the `rliable`.

D.5 STABILITY VERSUS PLASTICITY: INCREASING THE SIZE OF THE REPLAY BUFFER

By increasing the replay buffer size for world model learning for DreamerV2 + Plan2Explore we see that forgetting and average performance increase. However, the forward transfer simultaneously decreases, [Figure 12](#). Additionally, by inspecting the learning curves we notice that the harder exploration tasks are not learned as the replay buffer size increases. This is an instance of the stability-plasticity trade-off in continual learning the larger buffer size enables better remembering but simultaneously prevents quickly learning the new tasks.

D.6 INTERFERENCE

We observe interference explicitly using DreamerV2 with Plan2Explore on the `Minigrid FourRooms-v0` environment where we change the reward function or goal location from one task to another [Figure 13](#). We train for 1M steps on the first task, then 1M steps on the second task where the goal location has changed with all else the same. From the learning curves for each individual run, we can see how only one task will ever be solved with simultaneously poor performance in the second task (with the exception of one seed, in blue, which switches between solving and not solving each task).

D.7 DECREASING THE REPLAY BUFFER SIZE

We decrease the size of the replay buffer for DreamerV2 and variants to see how well it is able to manage environment reconstruction in the face of continual learning and decreasing replay buffer size. We consider the 3 task `Minigrid` continual learning problem and decrease the size of the replay buffer from 2×10^6 transitions which is used in [Section 5.1](#) to replay buffers of size $\{10^4, 10^5, 10^6\}$ transitions.

From the results in [Figure 14](#), we can see that DreamerV2 and its variants under-perform with small replay buffers of size 10^4 and 10^5 . DreamerV2 with Plan2Explore is unable to learn with such small replay buffers. DreamerV2 by itself is better at learning under small replay buffers. DreamerV2 with Plan2Explore learns to solve the difficult `DoorKey-9x9` problem only when it has a replay buffer size of 10^6 . We can also see that reservoir sampling helps against catastrophic forgetting for both Dreamer and DreamerV2 with Plan2Explore with replay buffer sizes of 10^5 and 10^6 .

D.8 ANALYSIS OF REPLAY BUFFER SAMPLING METHODS

We analyze the workings of the different minibatch construction sampling methods: random sampling, *us*, *rwd* and *50:50* sampling. For *50 : 50* sampling we also employ reservoir sampling (since in our experiments in [Section 5](#) we employ it to add plasticity to *rs*). All other sampling methods use a FIFO replay buffer. We plot histograms of the episode index

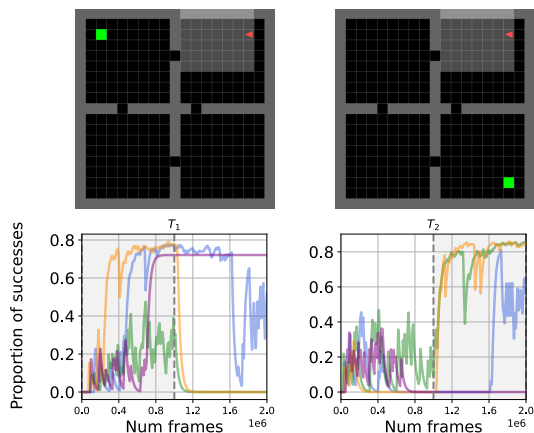


Figure 13: **Top**, 2 `FourRooms-v0` environment on with fixed agent start location, agent start orientation, and obstacles. Only the goal or reward function changes from one task to the next. **Bottom**, 4 separate success rate learning curves for different random seeds in different colours for DreamerV2 + p2e.

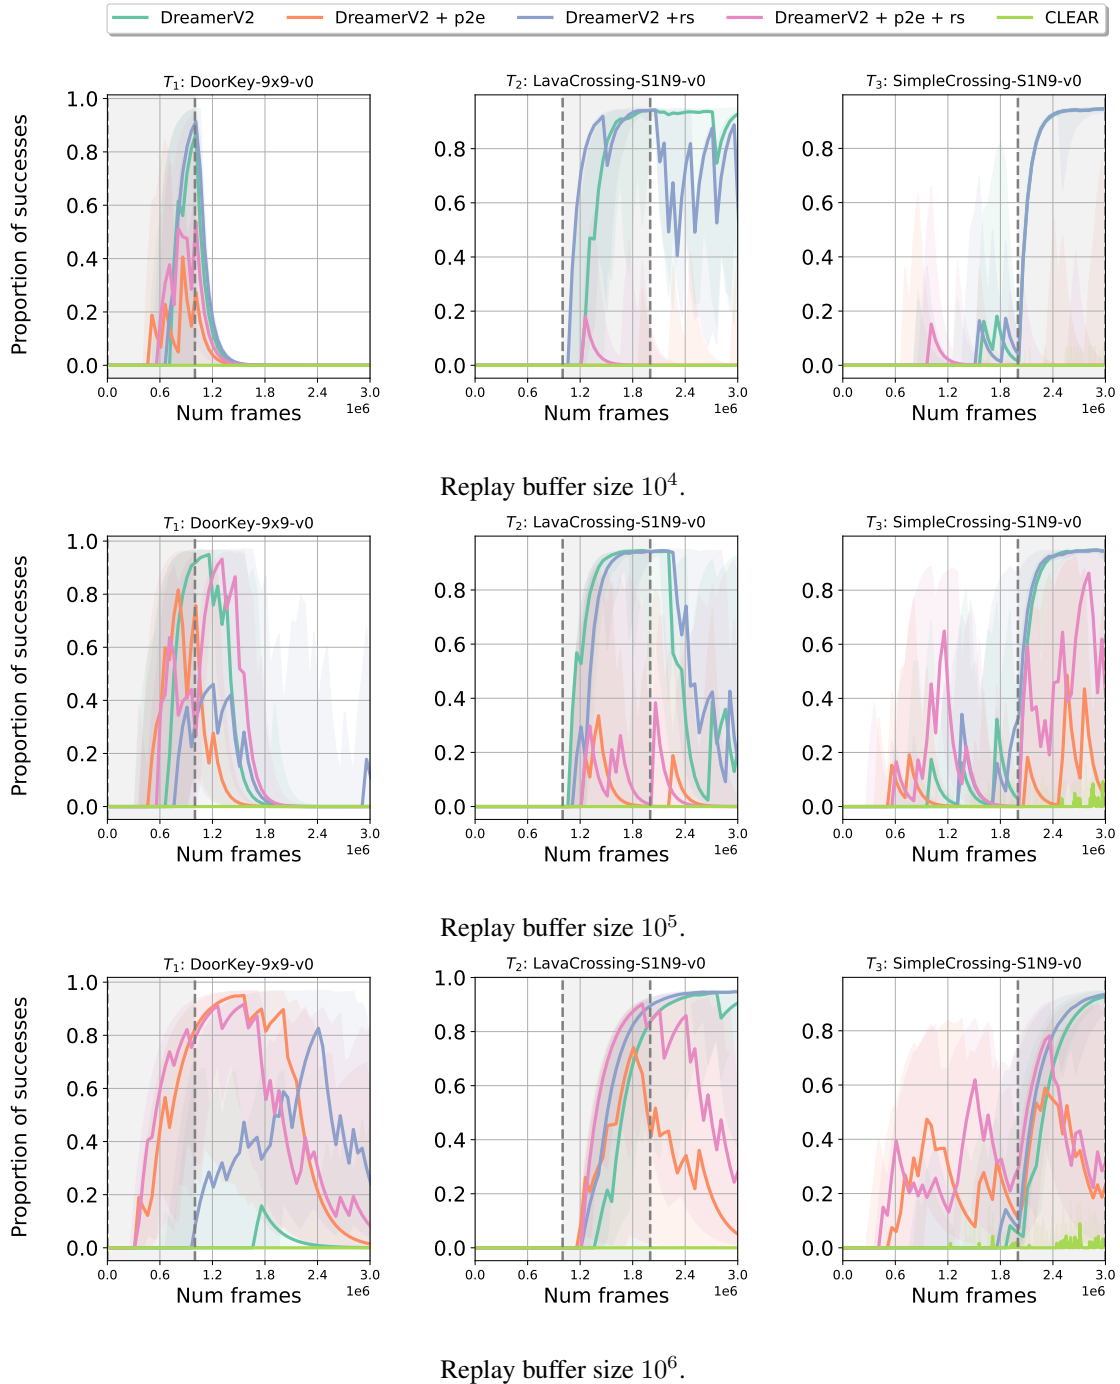


Figure 14: Learning curves for continual learning on 3 different Minigrad tasks with 1M environment interactions per task before changing task. The replay buffer is decreased from the experiments in Section 5.1 to 10^4 transitions in the top row, 10^5 in the middle row, and 10^6 in the bottom row. All runs are medians and interquartile ranges of over 10 different runs with different seeds.

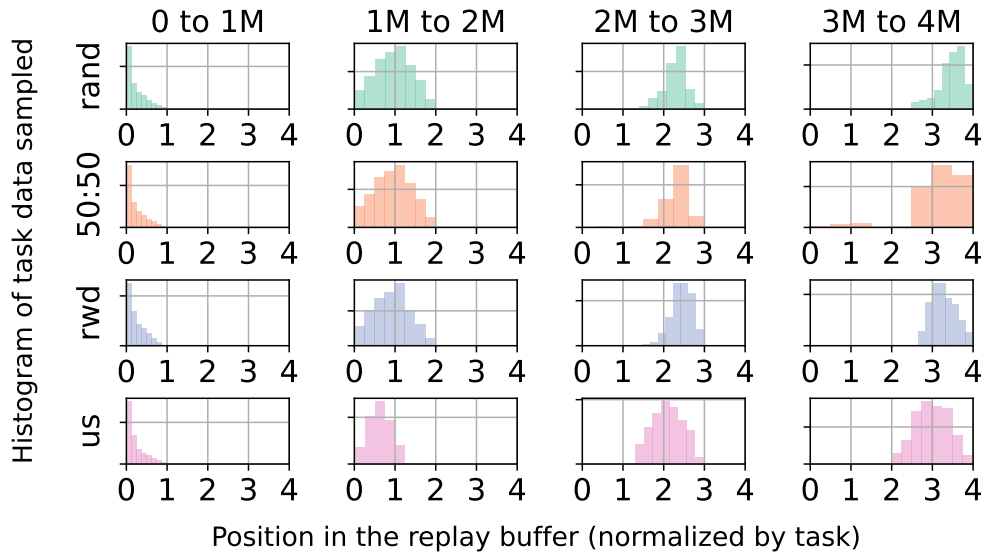


Figure 15: Distribution of episode indexes which are sampled to construct a minibatch for various sampling methods (rows) while the world-model is interacting with each task (columns).

which is sampled for mini-batch construction for world-model learning and initiating the Actor-Critic learning. We normalized the replay buffer episode indexes, which are sampled so that they lie in the range $[0, 4]$ to indicate which task the episode corresponds to. So all the episode indexes which were sampled while the agent interacted with a particular task are visualized in a distinct histogram (columns) for various sampling methods (rows) [Figure 15](#).

We can see from [Figure 15](#) that the sampling from the replay buffer is similar across all sampling methods for the first task. The distributions for random sampling are not uniform since only episodes which are of length > 50 are stored in the replay buffer, so the histograms' first row is not flat since certain indexes will not be sampled as a result. The distribution for *rwd* and *us* is similar to random sampling for the first task. This must mean that the uncertainty from Plan2Explore is quite uniform for the first task, perhaps even the uncertainty is larger for the earlier episodes in the replay buffer for the first task (row 4, column 1 [Figure 15](#)). As continual learning progresses we can see how the sampling from the replay buffer becomes more uneven in comparison to random sampling for the next tasks for *rwd* and *us*. In particular, we can see for time-steps 3M to 4M the detrimental effects of *rwd*, since only previous experience with high rewards are sampled when learning the final task causing a lack of plasticity for *rwd*. When looking at the distributions for 50 : 50 sampling we can see how they are more shifted to the right in comparison to random sampling, this is to be expected since we are explicitly biasing the sampling to more recent experience to make the world model learn recent tasks quicker since 50 : 50 is paired with *rs*.