

T-DGR: A TRAJECTORY-BASED DEEP GENERATIVE REPLAY METHOD FOR CONTINUAL LEARNING IN DECISION MAKING

William Yue

The University of Texas at Austin
william.yue@utexas.edu

Bo Liu

The University of Texas at Austin
bliu@cs.utexas.edu

Peter Stone

The University of Texas at Austin
Sony AI
pstone@cs.utexas.edu

ABSTRACT

Deep generative replay has emerged as a promising approach for continual learning in decision-making tasks. This approach addresses the problem of catastrophic forgetting by leveraging the generation of trajectories from previously encountered tasks to augment the current dataset. However, existing deep generative replay methods for continual learning rely on autoregressive models, which suffer from compounding errors in the generated trajectories. In this paper, we propose a simple, scalable, and non-autoregressive method for continual learning in decision-making tasks using a generative model that generates task samples conditioned on the trajectory timestep. We evaluate our method on Continual World benchmarks and find that our approach achieves state-of-the-art performance on the average success rate metric among continual learning methods.

1 INTRODUCTION

Continual learning, also known as lifelong learning, is a critical challenge in the advancement of general artificial intelligence, as it enables models to learn from a continuous stream of data encompassing various tasks, rather than having access to all data at once (Ring, 1994). However, a major challenge in continual learning is the phenomenon of catastrophic forgetting, where previously learned skills are lost when attempting to learn new tasks (McCloskey & Cohen, 1989).

To mitigate catastrophic forgetting, replay methods have been proposed, which involve saving data from previous tasks and replaying it to the learner during the learning of future tasks. This approach mimics how humans actively prevent forgetting by reviewing material for tests and replaying memories in dreams. However, storing data from previous tasks requires significant storage space and becomes computationally infeasible as the number of tasks increases.

In the field of cognitive neuroscience, the Complementary Learning Systems theory offers insights into how the human brain manages memory. This theory suggests that the brain employs two complementary learning systems: a fast-learning episodic system and a slow-learning semantic system (McClelland et al., 1995; Kumaran et al., 2016; McClelland et al., 1995). The hippocampus serves as the episodic system, responsible for storing specific memories of unique events, while the neocortex functions as the semantic system, extracting general knowledge from episodic memories and organizing it into abstract representations (O’Reilly & Norman, 2002).

Drawing inspiration from the human brain, deep generative replay (DGR) addresses the catastrophic forgetting issue in decision-making tasks by using a generative model as the hippocampus to generate trajectories from past tasks and replay them to the learner which acts as the neocortex (Figure 2) (Shin et al., 2017). The time-series nature of trajectories in decision-making tasks sets it apart from continual supervised learning, as each timestep of the trajectory requires sufficient replay. In supervised learning, the learner’s performance is not significantly affected if it performs poorly on a small subset of the data. However, in decision-making tasks, poor performance on any part of the trajectory can severely impact the overall performance. Therefore, it is crucial to generate state-action pairs that accurately represent the distribution found in trajectories. Furthermore, the high-dimensional distribution space of trajectories makes it computationally infeasible to generate complete trajectories all at once.

Existing DGR methods adopt either the generation of individual state observations i.i.d. without considering the temporal nature of trajectories or autoregressive trajectory generation. Autoregressive approaches generate the next state(s) in a trajectory by modeling the conditional probability of the next state(s) given the previously generated state(s). However, autoregressive methods suffer from compounding errors in the generated trajectories. On the other

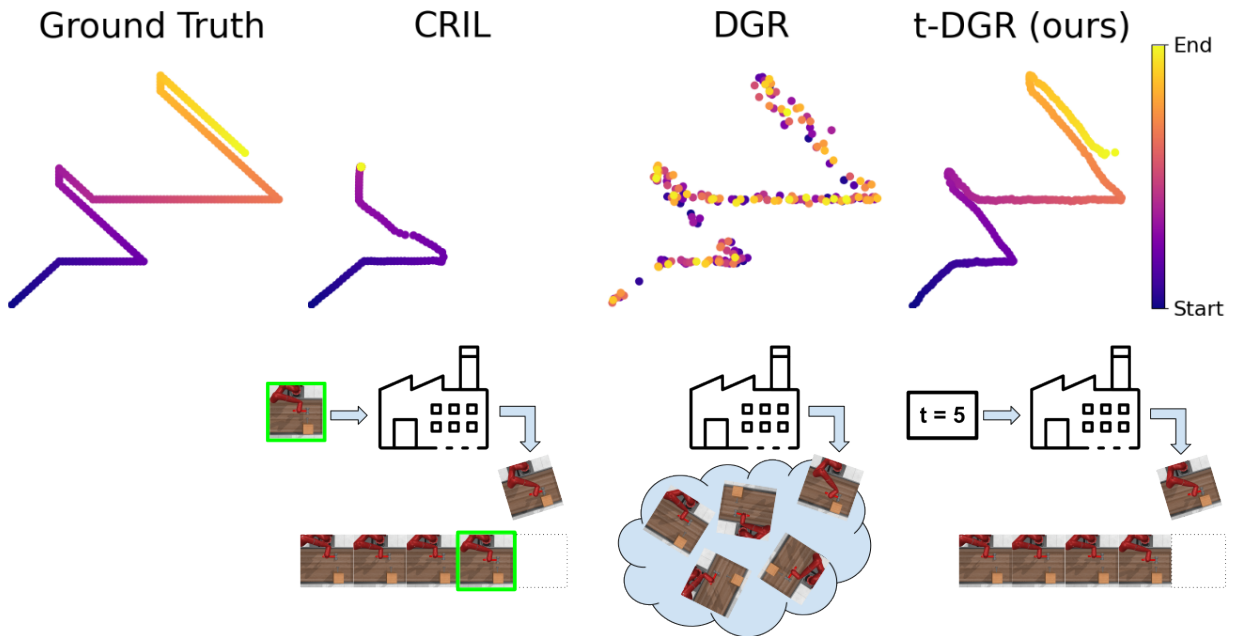


Figure 1: The first row presents a comparison of three generative methods for imitating an agent’s movement in a continuous 2D plane with Gaussian noise. The objective is to replicate the ground truth path, which transitions from darker to lighter colors. The autoregressive method (CRIL) encounters a challenge at the first sharp turn as nearby points move in opposing directions. Once the autoregressive method deviates off course, it never recovers and compromises the remaining trajectory. In contrast, sampling individual state observations i.i.d. without considering the temporal nature of trajectories (DGR) leads to a fragmented path with numerous gaps. Our proposed method t-DGR samples individual state observations conditioned on the trajectory timestep. By doing so, t-DGR successfully avoids the pitfalls of CRIL and DGR, ensuring a more accurate replication of the desired trajectory. The second row illustrates how each method generates trajectory data. CRIL generates the next state observation conditioned on the previous state observation. DGR, in contrast, does not attempt to generate a trajectory but generates individual state observations i.i.d. On the other hand, t-DGR generates state observations conditioned on the trajectory timestep.

hand, generating individual state observations i.i.d. leads to a higher sample complexity compared to generating entire trajectories, which becomes significant when replay time is limited (see Section 2.2).

To address the issues in current DGR methods, we propose a simple, scalable, and non-autoregressive trajectory-based DGR method. We define a generated trajectory as temporally coherent if the transitions from one state to the next appear realistic (refer to Section 3.4 for a formal definition). Given that imitation learning methods and many reinforcement learning methods are trained on single transition samples (e.g. state-action pairs), we do not require trajectories to exhibit temporal coherence. Instead, our focus is on ensuring an equal number of samples generated at each timestep of the trajectory to accurately represent the distribution found in trajectories. To achieve equal sample coverage at each timestep, we train our generator to produce state observations conditioned on the trajectory timestep, and then sample from the generator conditioned on each timestep of the trajectory. The intuition behind our method is illustrated in Figure 1.

To evaluate the effectiveness of our proposed method, t-DGR, we conducted experiments on the Continual World benchmarks CW10 and CW20 (Wolczyk et al., 2021) using imitation learning. Our results indicate that t-DGR achieves state-of-the-art performance in terms of average success rate when compared to other top continual learning methods.

2 RELATED WORK

This section provides an overview of existing continual learning methods, with a particular focus on pseudo-rehearsal methods.

2.1 CONTINUAL LEARNING IN THE REAL WORLD

As the field of continual learning continues to grow, there is an increasing emphasis on developing methods that can be effectively applied in real-world scenarios (Wang et al., 2024; Aljundi et al., 2019; Bang et al., 2021; Hsu et al., 2019; van de Ven & Tolia, 2019). The concept of “General Continual Learning” was introduced by Buzzega et al. (2020) to address certain properties of the real world that are often overlooked or ignored by existing continual learning methods. Specifically, two important properties, bounded memory and blurry task boundaries, are emphasized in this work. Bounded memory refers to the requirement that the memory footprint of a continual learning method should remain bounded throughout the entire lifespan of the learning agent. This property is crucial to ensure practicality and efficiency in real-world scenarios. Additionally, blurry task boundaries highlight the challenge of training on tasks that are intertwined, without clear delineation of when one task ends and another begins. Many existing methods fail to account for this characteristic, which is common in real-world learning scenarios. While there are other significant properties associated with continual learning in the real world, this study focuses on the often-neglected aspects of bounded memory and blurry task boundaries. By addressing these properties, we aim to develop methods that are more robust and applicable in practical settings.

2.2 CONTINUAL LEARNING METHODS

Continual learning methods for decision-making tasks can be categorized into three main categories.

Regularization Regularization methods in continual learning focus on incorporating constraints during model training to promote the retention of past knowledge. One simple approach is to include an L_2 penalty in the loss function. Elastic Weight Consolidation (EWC) builds upon this idea by assigning weights to parameters based on their importance for previous tasks using the Fisher information matrix (Kirkpatrick et al., 2017). MAS measures the sensitivity of parameter changes on the model’s output, prioritizing the retention of parameters with a larger effect (Aljundi et al., 2018). VCL leverages variational inference to minimize the Kullback-Leibler divergence between the current and prior parameter distributions (Nguyen et al., 2018). Progress and Compress learns new tasks using a separate model and subsequently distills this knowledge into the main model while safeguarding the previously acquired knowledge (Schwarz et al., 2018). However, regularization methods may struggle with blurry task boundaries as they rely on knowledge of task endpoints to apply regularization techniques effectively. In our experiments, EWC was chosen as the representative regularization method based on its performance in the original Continual World experiments (Wołczyk et al., 2021).

Architecture-based Methods Architecture-based methods aim to maintain distinct sets of parameters for each task, ensuring that future learning does not interfere with the knowledge acquired from previous tasks. Packnet (Mallya & Lazebnik, 2018), UCL (Ahn et al., 2019), and AGS-CL (Jung et al., 2020) all safeguard previous task information in a neural network by identifying important parameters and freeing up less important parameters for future learning. Identification of important parameters can be done through iterative pruning (Packnet), parameter uncertainty (UCL), and activation value (AGS-CL). However, a drawback of parameter isolation methods is that each task requires its own set of parameters, which may eventually exhaust the available parameters for new tasks and necessitate a dynamically expanding network without bounded memory (Yoon et al., 2018). Additionally, parameter isolation methods require training on a single task at a time to prune and isolate parameters, preventing concurrent learning from multiple interwoven tasks. In our experiments, PackNet was selected as the representative architecture-based method based on its performance in the original Continual World experiments (Wołczyk et al., 2021).

Pseudo-rehearsal Methods Pseudo-rehearsal methods mitigate the forgetting of previous tasks by generating synthetic samples from past tasks and replaying them to the learner. Deep generative replay (DGR) (Figure 2) utilizes a generative model, such as generative adversarial networks (Goodfellow et al., 2020), variational autoencoders (Kingma & Welling), or diffusion models (Ho et al., 2020; Sohl-Dickstein et al., 2015), to generate the synthetic samples. Originally, deep generative replay was proposed to address continual supervised learning problems, where the generator only needed to generate single data point samples (Shin et al., 2017). However, in decision-making tasks, expert demonstrations consist of trajectories (time-series) with a significantly higher-dimensional distribution space.

One existing DGR method generates individual state observations i.i.d. instead of entire trajectories. However, this approach leads to a higher sample complexity compared to generating entire trajectories. The sample complexity of generating enough individual state observations i.i.d. to cover every portion of the trajectory m times can be described using the Double Dixie Cup problem (Newman, 1960). For trajectories of length n , it takes an average of $\Theta(n \log n + mn \log \log n)$ i.i.d. samples to ensure at least m samples for each timestep. In scenarios with limited replay time (small m) and long trajectories (large n) the sample complexity can be approximated as $\Theta(n \log n)$ using




Figure 2: The deep generative replay paradigm. The algorithm learns to generate trajectories from past tasks to augment real trajectories from the current task in order to mitigate catastrophic forgetting. Both the generator and policy model are updated with this augmented dataset.

the Coupon Collector’s problem (Myers & Wilf, 2006). The additional $\Theta(\log n)$ factor reduces the likelihood of achieving complete sample coverage of the trajectory when the number of replays or replay time is limited, especially considering the computationally expensive nature of current generative methods. Furthermore, there is a risk that the generator assigns different probabilities to each timestep of the trajectory, leading to a selective focus on certain timesteps rather than equal representation across the trajectory.

Another existing DGR method is autoregressive trajectory generation. In the existing autoregressive method, CRIL, a generator is used to generate samples of the initial state, and a dynamics model predicts the next state based on the current state and action (Gao et al., 2021). However, even with a dynamics model accuracy of 99% and a 1% probability of deviating from the desired trajectory, the probability of an autoregressively generated trajectory going off course is $1 - 0.99^n$, where n denotes the trajectory length. With a trajectory length of $n = 200$ (as used in our experiments), the probability of an autoregressively generated trajectory going off course is $1 - 0.99^{200} = 0.87$. This example demonstrates how the issue of compounding error leads to a high probability of failure, even with a highly accurate dynamics model.

In our experiments, t-DGR is evaluated against all existing trajectory generation methods in pseudo-rehearsal approaches to assess how well t-DGR addresses the limitations of those methods.

3 BACKGROUND

This section introduces notation and the formulation of the continual imitation learning problem that we use in this paper. Additionally, we provide a concise overview of diffusion probabilistic models used in our generative model implementation.

3.1 IMITATION LEARNING

Imitation learning algorithms aim to learn a policy π_θ parameterized by θ by imitating a set of expert demonstrations $D = \{\tau_i\}_{i=1\dots M}$. Each trajectory τ_i consists of a sequence of state-action pairs $\{(s_j, a_j)\}_{j=1\dots|\tau_i|}$ where $|\tau_i|$ is the length of the trajectory. Each trajectory comes from a task \mathcal{T} which is a Markov decision process that can be represented as a tuple $\langle S, A, T, \rho_0 \rangle$ with state space S , action space A , transition dynamics $T : S \times A \times S \rightarrow [0, 1]$, and initial state distribution ρ_0 . Various algorithms exist for imitation learning, including behavioral cloning, GAIL (Ho & Ermon, 2016), and inverse reinforcement learning (Ng & Russell, 2000). In this work, we use behavioral cloning where the objective can be formulated as minimizing the loss function:

$$\mathcal{L}(\theta) = \mathbb{E}_{s,a \sim D} \left[\left\| \pi_\theta(s) - a \right\|_2^2 \right] \quad (1)$$

where the state and action spaces are continuous.

3.2 CONTINUAL IMITATION LEARNING

In the basic formulation most common in the field today, continual imitation learning involves sequentially solving multiple tasks $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_N$. When solving for task \mathcal{T}_i , the learner only gets data from task \mathcal{T}_i and can not access data for any other task. In a more general scenario, certain tasks may have overlapping boundaries, allowing the learner to encounter training data from multiple tasks during certain phases of training. The learner receives a continuous stream of training data in the form of trajectories $\tau_1, \tau_2, \tau_3, \dots$ from the environment, where each trajectory τ corresponds to one of the N tasks. However, the learner can only access a limited contiguous portion of this stream at any given time.

Let s_i be the success rate of task \mathcal{T}_i after training on all N tasks. The continual imitation learning objective is defined as maximizing the average success rate over all tasks:

$$S = \frac{1}{N} \sum_{i=1}^N s_i \quad (2)$$

The primary issue that arises from the continual learning problem formulation is the problem of catastrophic forgetting where previously learned skills are forgotten when training on a new task.

3.3 DIFFUSION PROBABILISTIC MODELS

Diffusion probabilistic models (Ho et al., 2020; Sohl-Dickstein et al., 2015) generate data through a learned reverse denoising diffusion process $p_\theta(x_{t-1} | x_t)$. The forward diffusion process $q(x_t | x_{t-1})$ gradually adds Gaussian noise to an input x_0 at each time step t , ultimately resulting in pure noise x_T at $t = T$. The forward diffusion process is defined as:

$$q(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}) \quad (3)$$

where $0 < \beta_t < 1$ is defined by a known variance schedule. In our implementation, we adopted the cosine schedule proposed by Nichol & Dhariwal (2021). For a sufficiently large time horizon T and a well-behaved variance schedule, x_T approximates an isotropic Gaussian distribution. If we had the reverse diffusion process $p(x_{t-1} | x_t)$, we could sample $x_T \sim \mathcal{N}(0, \mathbf{I})$ and obtain a sample from $q(x_0)$ by denoising x_T with $p(x_{t-1} | x_t)$. However, computing $p(x_{t-1} | x_t)$ is intractable as it necessitates knowledge of the distribution of all possible x_t . Instead, we approximate $p(x_{t-1} | x_t)$ using a neural network:

$$p_\theta(x_{t-1} | x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma(x_t, t)) \quad (4)$$

Since q and p_θ can be viewed as a variational auto-encoder (Kingma & Welling), we can use the variational lower bound to minimize the negative log-likelihood of the reverse process. We can express $\mu_\theta(x_t, t)$ from Equation 4 as:

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right) \quad (5)$$

where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=0}^t \alpha_s$. The training loss can then be defined as:

$$\mathcal{L}(\theta) = \mathbb{E}_{x_0, t, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2] \quad (6)$$

Note that the timesteps t in the diffusion process differ from the trajectory timesteps t . Henceforth, we will refer only to the trajectory timesteps t .

3.4 NOTATION

Deep generative replay involves training two models: a generator G_γ parameterized by γ and a learner π_θ parameterized by θ . We define $G_\gamma^{(i)}$ as the generator trained on tasks $\mathcal{T}_1 \dots \mathcal{T}_i$ and capable of generating data samples from tasks $\mathcal{T}_1 \dots \mathcal{T}_i$. Similarly, $\pi_\theta^{(i)}$ represents the learner trained on tasks $\mathcal{T}_1 \dots \mathcal{T}_i$ and able to solve tasks $\mathcal{T}_1 \dots \mathcal{T}_i$.

A sequence of state observations $(s_1, s_2, \dots, s_{n-1}, s_n)$ is **temporally coherent** if $\forall 1 \leq i < n, \exists a \in A : T(s_i, a, s_{i+1}) > \varepsilon$, where $0 < \varepsilon < 1$ is a small constant representing a threshold for negligible probabilities.

4 METHOD

Our proposed method, t-DGR, tackles the challenge of generating long trajectories by training a generator, denoted as $G_\gamma(j)$, which is conditioned on the trajectory timestep j to generate state observations. Pseudocode for t-DGR is provided as Algorithm 1. The algorithm begins by initializing the task index, replay ratio, generator model, learner model, and learning rates (Line 1). The replay ratio, denoted as $0 \leq r < 1$, determines the percentage of training samples seen by the learner that are generated. Upon receiving training data from the environment, t-DGR calculates the number of trajectories to generate based on the replay ratio r (Lines 4-5). The variable L (Line 7) represents the maximum length of trajectories observed so far.

To generate a trajectory τ of length L , t-DGR iterates over each timestep $1 \leq j \leq L$ (Line 9). At each timestep, t-DGR generates the j -th state observation of the trajectory using the previous generator $G_\gamma^{(t-1)}$ conditioned on timestep j

Algorithm 1 Trajectory-based Deep Generative Replay (t-DGR)

```

1: Initialize task index  $t = 0$ , replay ratio  $r$ , generator  $G_\gamma^{(0)}$ , learner  $\pi_\theta^{(0)}$ , and learning rates  $\lambda_\gamma, \lambda_\theta$ .
2: while new task available do
3:    $t \leftarrow t + 1$ 
4:   Initialize dataset  $D$  with trajectories from task  $t$ .
5:    $n \leftarrow \frac{r*|D|}{1-r}$  ▷ number of trajectories to generate
6:   for  $i = 1$  to  $n$  do
7:      $L \leftarrow$  maximum trajectory length
8:      $\tau \leftarrow \emptyset$  ▷ initialize trajectory of length  $L$ 
9:     for  $j = 1$  to  $L$  do
10:       $S \leftarrow G_\gamma^{(t-1)}(j)$  ▷ generate states
11:       $A \leftarrow \pi_\theta^{(t-1)}(S)$  ▷ label with actions
12:       $\tau_j \leftarrow (S, A)$  ▷ add to trajectory
13:    end for
14:     $D \leftarrow D \cup \tau$  ▷ add generated trajectory to  $D$ 
15:  end for
16:  Update generator and learner using  $D$ 
17:   $\gamma^{(t)} \leftarrow \gamma^{(t-1)} - \lambda_\gamma \nabla_\gamma \mathcal{L}_{G^{(t-1)}}(\gamma^{(t-1)})$ 
18:   $\theta^{(t)} \leftarrow \theta^{(t-1)} - \lambda_\theta \nabla_\theta \mathcal{L}_{\pi^{(t-1)}}(\theta^{(t-1)})$ 
19: end while

```

(Line 10), and then labels it with an action using the previous policy $\pi_\theta^{(t-1)}$ (Line 11). After generating all timesteps in the trajectory τ , t-DGR adds it to the existing training dataset (Line 14). Note that the generated state observations within a trajectory do not have temporal coherence, as each state observation is generated independently of other timesteps. This approach is acceptable since our learner is trained on state-action pairs rather than full trajectories. However, unlike generating state observations i.i.d., our method ensures equal coverage of every timestep during the generative process, significantly reducing sample complexity.

Once t-DGR has augmented the training samples from the environment with our generated training samples, t-DGR employs backpropagation to update both the generator and learner using the augmented dataset (Lines 16-18). The t-DGR algorithm continues this process of generative replay throughout the agent’s lifetime, which can be infinite (Line 2). Although we perform the generative process of t-DGR at task boundaries for ease of understanding, no part of t-DGR is dependent on clear task boundaries.

Architecture We employ a U-net (Ronneberger et al., 2015) trained on the loss specified in Equation 6 to implement the generative diffusion model G_γ . Since we utilize proprioceptive observations in our experiments, π_θ is implemented with a multi-layer perceptron trained on the loss specified in Equation 1.

5 EXPERIMENTS

In this section, we outline the experimental setup and performance metrics employed to compare t-DGR with representative methods, followed by an analysis of experimental results across different benchmarks and performance metrics.

5.1 EXPERIMENTAL SETUP

We evaluate our method on the Continual World benchmarks CW10 and CW20 (Wolczyk et al., 2021), along with our own variant of CW10 called BB10 that evaluates the ability of methods to handle blurry task boundaries. CW10 consists of a sequence of 10 Meta-World (Yu et al., 2020) tasks, where each task involves a Sawyer arm manipulating one or two objects in the Mujoco physics simulator. For computational efficiency, we provide the agents with proprioceptive observations. Notably, the observation and action spaces are continuous and remain consistent across all tasks. CW20 is an extension of CW10 with the tasks repeated twice. To our knowledge, Continual World is the only standard continual learning benchmark for decision-making tasks. BB10 gives data to the learner in 10 sequential buckets B_1, \dots, B_{10} . Data from task \mathcal{T}_i from CW10 is split evenly between buckets B_{i-1} , B_i , and B_{i+1} , except for

the first and last task. Task \mathcal{T}_1 is evenly split between buckets B_0 and B_1 , and task \mathcal{T}_{10} is evenly split between buckets B_9 and B_{10} .

In order to bound memory usage, task conditioning should utilize fixed-size embeddings like natural language task embeddings, rather than maintaining a separate final neural network layer for each individual task. For simplicity, we condition the model on the task with a one-hot vector as a proxy for natural language prompt embeddings. For BB10, the model is still conditioned on the underlying task rather than the bucket. Additionally, we do not allow separate biases for each task, as originally done in EWC (Kirkpatrick et al., 2017). Expert demonstrations for training are acquired by gathering 100 trajectories per task using hand-designed policies from Meta-World, with each trajectory limited to a maximum of 200 steps. Importantly, the learner model remains consistent across different methods and benchmark evaluations. Moreover, we maintain a consistent replay ratio of $r = 0.9$ across all pseudo-rehearsal methods.

We estimated the success rate S of a model by running each task 100 times. The metrics for each method were computed using 5 seeds to create a 90% confidence interval. Further experimental details, such as hyperparameters, model architecture, random seeds, and computational resources, are included in the appendix. This standardization enables a fair and comprehensive comparison of our proposed approach with other existing methods.

5.2 METRICS

We evaluate our models using three metrics proposed by the Continual World benchmark (Wołczyk et al., 2021), with the average success rate being the primary metric. Although the forward transfer and forgetting metrics are not well-defined in a setting with blurry task boundaries, they are informative within the context of Continual World benchmarks. As a reminder from Section 3.2, let N denote the number of tasks, and s_i represent the success rate of the learner on task \mathcal{T}_i . Additionally, let $s_i(t)$ denote the success rate of the learner on task \mathcal{T}_i after training on tasks \mathcal{T}_1 to \mathcal{T}_t .

Average Success Rate The average success rate, as given by Equation 2, serves as the primary evaluation metric for continual learning methods.

Average Forward Transfer We introduce a slightly modified metric for forward transfer that applies to a broader range of continual learning problems beyond just continual reinforcement learning in the Continual World benchmark. Let s_i^{ref} represent the reference performance of a single-task experiment on task \mathcal{T}_i . The forward transfer metric FT_i is computed as follows:

$$FT_i = \frac{D_i - D_i^{\text{ref}}}{1 - D_i^{\text{ref}}} \quad D_i = \frac{s_i(i) + s_i(i-1)}{2} \quad D_i^{\text{ref}} = \frac{s_i^{\text{ref}}}{2}$$

The expressions for D_i and D_i^{ref} serve as approximations of the integral of task \mathcal{T}_i performance with respect to the training duration for task \mathcal{T}_i . The average forward transfer FT is then defined as the mean forward transfer over all tasks, calculated as $FT = \frac{1}{N} \sum_{i=1}^N FT_i$.

Average Forgetting We measure forgetting using the metric F_i , which represents the amount of forgetting for task i after all training has concluded. F_i is defined as the difference between the success rate on task \mathcal{T}_i immediately after training and the success rate on task \mathcal{T}_i at the end of training.

$$F_i = s_i(i) - s_i(N)$$

The average forgetting F is then computed as the mean forgetting over all tasks, given by $F = \frac{1}{N} \sum_{i=1}^N F_i$.

5.3 BASELINES

We compare the following methods on the Continual World benchmark using average success rate as the primary evaluation metric. Representative methods were chosen based on their success in the original Continual World experiments, while DGR-based methods were selected to evaluate whether t-DGR addresses the limitations of existing pseudo-rehearsal methods.

- **Finetune:** The policy is trained only on data from the current task.
- **Multitask:** The policy is trained on data from all tasks simultaneously.

- **oEWC (Schwarz et al., 2018)**: A variation of EWC known as online Elastic Weight Consolidation (oEWC) bounds the memory of EWC by employing a single penalty term for the previous model instead of individual penalty terms for each task. This baseline is the representative regularization-based method.
- **PackNet (Mallya & Lazebnik, 2018)**: This baseline is the representative parameter isolation method. PackNet safeguards previous task information in a neural network by iteratively pruning, freezing, and retraining parts of the network.
- **DGR (Shin et al., 2017)**: This baseline is a deep generative replay method that only generates individual state observations i.i.d. and not entire trajectories.
- **CRIL (Gao et al., 2021)**: This baseline is a deep generative replay method that trains a policy along with a start state generator and a dynamics model that predicts the next state given the current state and action. Trajectories are generated by using the dynamics model and policy to autoregressively generate next states from a start state.
- **t-DGR**: Our proposed method.

Due to the inability of oEWC and PackNet to handle blurry task boundaries, we made several adjustments for CW20 and BB10. Since PackNet cannot continue training parameters for a task once they have been fixed, we treated the second repetition of tasks in CW20 as distinct from the first iteration, resulting in PackNet being evaluated with $N = 20$, while the other methods were evaluated with $N = 10$. As for BB10 and its blurry task boundaries, the best approach we could adopt with oEWC and PackNet was to apply their regularization techniques at regular training intervals rather than strictly at task boundaries. During evaluation, all tasks were assessed using the last fixed set of parameters in the case of PackNet.

5.4 ABLATIONS

To evaluate the effect of generator architecture on pseudo-rehearsal methods, we included an ablation where the diffusion models in pseudo-rehearsal methods are replaced with Wasserstein generative adversarial networks (GAN) with gradient penalty (Gulrajani et al., 2017). To evaluate the effectiveness of diffusion models at generative replay, we evaluated the quality of generated samples for past tasks at all stages of continual learning. Since we utilize proprioceptive observations rather than RGB camera observations, there is no practical way to qualitatively evaluate the generated samples. The proprioceptive observations are 39-dimensional, comprising 3D positions and quaternions of objects, and the degree to which the robot gripper is open (Yu et al., 2020). These raw numbers cannot be qualitatively evaluated like RGB images. Instead, we use the average diffusion loss as a quantitative proxy metric for generation quality. When the generator model is learning task i , we compute the average diffusion loss for data samples from tasks 1 to $i - 1$. We then compare this average diffusion loss to the average diffusion loss when the generator model first learned task i . The difference between the two provides an estimate of how much generative replay has degraded the generation quality of previous tasks.

5.5 DISCUSSION

t-DGR emerges as the leading method, demonstrating the highest success rate on CW10 (Table 1a), CW20 (Table 1c), and BB10 (Table 1b). Notably, PackNet’s performance on the second iteration of tasks in CW20 diminishes, highlighting its limited capacity for continually accommodating new tasks. This limitation underscores the fact that PackNet falls short of being a true lifelong learner, as it necessitates prior knowledge of the task count for appropriate parameter capacity allocation. On the contrary, pseudo-rehearsal methods, such as t-DGR, exhibit improved performance with the second iteration of tasks in CW20 due to an increased replay time. These findings emphasize the ability of DGR methods to effectively leverage past knowledge, as evidenced by their superior forward transfer in both CW10 and CW20.

BB10 (Table 1b) demonstrates that pseudo-rehearsal methods are mostly unaffected by blurry task boundaries, whereas PackNet’s success rate experiences a significant drop-off. This discrepancy arises from the fact that PackNet’s regularization technique does not work effectively with less clearly defined task boundaries.

In our experiments across CW10, CW20, and BB10, we observed that diffusion models outperform GANs as the generator for pseudo-rehearsal methods. We hypothesize that the distributional shifts in continual learning exacerbate instability issues with GAN training (Salimans et al., 2016; Brock et al., 2018; Miyato et al., 2018). The motivation behind CRIL as stated in the paper (Gao et al., 2021) is to alleviate the burden of trajectory generation from the generator by transferring some of the generation complexity to a dynamics model. Our findings support this motivation when working with less capable generators (Dhariwal & Nichol, 2021), as reducing the generator’s burden

(a) CW10				(b) BB10	
Method	Success Rate \uparrow	FT \uparrow	Forgetting \downarrow	Method	Success Rate \uparrow
Finetune	16.4 \pm 6.4	-3.0 \pm 6.0	78.8 \pm 7.6	Finetune	21.7 \pm 2.6
Multitask	97.0 \pm 1.0	N/A	N/A	Multitask	97.0 \pm 1.0
oEWC	18.6 \pm 5.3	-6.3 \pm 5.7	74.1 \pm 6.1	oEWC	21.8 \pm 1.7
PackNet	81.4 \pm 3.7	-14.8 \pm 7.8	-0.1 \pm 1.2	PackNet	26.9 \pm 5.6
DGR (gan)	21.2 \pm 3.7	-2.6 \pm 3.3	74.4 \pm 4.7	DGR (gan)	35.6 \pm 4.3
CRIL (gan)	24.0 \pm 4.4	0.3 \pm 4.5	72.0 \pm 5.4	CRIL (gan)	37.5 \pm 4.7
t-DGR (gan)	17.4 \pm 4.4	0.1 \pm 3.1	79.8 \pm 4.2	t-DGR (gan)	28.3 \pm 3.9
DGR (diffusion)	75.0 \pm 5.8	-4.3 \pm 5.1	17.8 \pm 4.1	DGR (diffusion)	75.3 \pm 4.4
CRIL (diffusion)	28.4 \pm 10.6	-1.1 \pm 2.8	68.6 \pm 10.4	CRIL (diffusion)	53.5 \pm 5.5
t-DGR (diffusion)	81.9 \pm 3.3	-0.3 \pm 4.9	14.4 \pm 2.5	t-DGR (diffusion)	81.7 \pm 4.0

(c) CW20				(d) Replay Ratio		
Method	Success Rate \uparrow	FT \uparrow	Forgetting \downarrow	Ratio	t-DGR	DGR
Finetune	14.2 \pm 4.0	-0.5 \pm 3.0	82.2 \pm 5.6	0.5	63.2 \pm 2.6	52.8 \pm 2.9
Multitask	97.0 \pm 1.0	N/A	N/A	0.6	66.3 \pm 4.4	56.9 \pm 4.5
oEWC	19.4 \pm 5.3	-2.8 \pm 4.1	75.2 \pm 7.5	0.7	70.8 \pm 4.1	62.5 \pm 3.6
PackNet	74.1 \pm 4.1	-20.4 \pm 3.4	-0.2 \pm 0.9	0.8	75.0 \pm 6.9	69.2 \pm 4.9
DGR (gan)	19.0 \pm 2.1	0.8 \pm 2.3	78.6 \pm 2.3	0.9	81.9 \pm 3.3	75.0 \pm 5.8
CRIL (gan)	26.7 \pm 6.5	-0.6 \pm 0.7	72.1 \pm 6.2			
t-DGR (gan)	20.6 \pm 6.1	0.7 \pm 3.4	76.3 \pm 5.3			
DGR (diffusion)	74.1 \pm 4.1	18.9 \pm 2.9	23.3 \pm 3.3			
CRIL (diffusion)	50.8 \pm 4.4	4.4 \pm 4.9	46.1 \pm 5.4			
t-DGR (diffusion)	83.9 \pm 3.0	30.6 \pm 4.5	14.6 \pm 2.9			

Table 1: Tables (a), (b), and (c) present the results for Continual World 10, Blurry Boundaries 10, and Continual World 20, respectively. The tables display the average success rate, forward transfer, and forgetting (if applicable) with 90% confidence intervals using 5 random seeds. An up arrow indicates that higher values are better and a down arrow indicates that smaller values are better. Table (d) compares the impact of replay amount on the average success rate of t-DGR and DGR on CW10 with 90% confidence intervals obtained using 5 random seeds. The best results are highlighted in bold.

consistently enhances performance of GAN-based pseudo-rehearsal methods across all benchmarks. Among pseudo-rehearsal techniques, CRIL places the least demand on its generator, which only needs to produce the initial frame of a trajectory. In contrast, DGR requires generation of every frame, and t-DGR, the most demanding, requires generation of every frame along with the handling of trajectory timestep conditioning. Results from our GAN-based pseudo-rehearsal experiments on CW10, CW20, and BB10 indicate that CRIL outperforms DGR, which in turn outperforms t-DGR. However, experiments using a more capable diffusion generator (Dhariwal & Nichol, 2021) reveal a reversal in performance ranking among these methods, suggesting that more capable generators diminish the need to offload generation complexity from the main generator.

The diminishing performance gap between DGR and t-DGR as the replay ratio increases in Table 1d indicates that a higher replay ratio reduces the likelihood of any portion of the trajectory being insufficiently covered when sampling individual state observations i.i.d., thereby contributing to improved performance. This trend supports the theoretical sample complexity of DGR derived in Section 2.2, as $\Theta(n \log n + mn \log \log n)$ closely approximates the sample complexity of t-DGR, $\Theta(mn)$, when the replay amount $m \rightarrow \infty$. However, while DGR can achieve comparable performance to t-DGR with a high replay ratio, the availability of extensive replay time is often limited in many real-world applications.

Recent studies have indicated that applying generative replay to diffusion models for image data leads to a collapse of generation quality due to compounding errors in the generated synthetic data (Zajac et al., 2023; Masip et al., 2023; Smith et al., 2023). However, our experiments reveal that when generating lower dimensional proprioceptive data, the diffusion model is able to maintain generation quality. Although Figure 3 shows that generative replay degrades generation quality of past tasks as the diffusion model learns new tasks, the success rates in Table 1 suggest that this degradation in generation quality is not severe enough to impact the learner’s performance. Notably, generation quality for past tasks appears to deteriorate in tasks 5 and 7 but improves in the subsequent task, suggesting that the error compounding in diffusion models during generative replay might not be as severe as previously assumed, particularly for lower-dimensional data.

Overall, t-DGR exhibits promising results, outperforming other methods in terms of success rate in all evaluations. Notably, t-DGR achieves a significant improvement over existing pseudo-rehearsal methods on CW20 using a Welch t-test with a significance level of $p\text{-value} = 0.005$. Its ability to handle blurry task boundaries, leverage past knowledge, and make the most of replay opportunities position it as a state-of-the-art method for continual lifelong learning in decision-making.




Figure 3: This table illustrates the ability of the diffusion model in t-DGR to generate past data as it continues to learn additional tasks in CW10 through generative replay. The line plot for task i plots the average diffusion loss of the diffusion model in future tasks on task i data. The loss is an L1 version of the diffusion training loss in Equation 6.

6 CONCLUSION

In conclusion, we have introduced t-DGR, a novel method for continual learning in decision-making tasks, which has demonstrated state-of-the-art performance on the Continual World benchmarks. Our approach stands out due to its simplicity, scalability, and non-autoregressive nature, positioning it as a solid foundation for future research in this domain.

Importantly, t-DGR takes into account essential properties of the real world, including bounded memory and blurry task boundaries. These considerations ensure that our method remains applicable and effective in real-world scenarios, enabling its potential integration into practical applications.

Looking ahead, one potential avenue for future research is the refinement of the replay mechanism employed in t-DGR. Rather than assigning equal weight to all past trajectories, a more selective approach could be explored. By prioritizing certain memories over others and strategically determining when to replay memories to the learner, akin to human learning processes, we could potentially enhance the performance and adaptability of our method.

ACKNOWLEDGEMENTS

This work has taken place in the Learning Agents Research Group (LARG) at the Artificial Intelligence Laboratory, The University of Texas at Austin. LARG research is supported in part by the National Science Foundation (FAIN-2019844, NRT-2125858), the Office of Naval Research (N00014-18-2243), Army Research Office (E2061621), Bosch, Lockheed Martin, and Good Systems, a research grand challenge at the University of Texas at Austin. The views and conclusions contained in this document are those of the authors alone. Peter Stone serves as the Executive Director of Sony AI America and receives financial compensation for this work. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research.

REFERENCES

- Hongjoon Ahn, Sungmin Cha, Donggyu Lee, and Taesup Moon. Uncertainty-based continual learning with adaptive regularization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/2c3ddf4bf13852db711dd1901fb517fa-Paper.pdf.
- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Computer Vision—ECCV 2018: 15th European Conference, Munich, Germany, September 8–14, 2018, Proceedings, Part III 15*, pp. 144–161. Springer, 2018.
- Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/e562cd9c0768d5464b64cf61da7fc6bb-Paper.pdf.
- Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8218–8227, 2021.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- 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.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- Chongkai Gao, Haichuan Gao, Shangqi Guo, Tianren Zhang, and Feng Chen. Cril: Continual robot imitation learning via generative and prediction model. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6747–6754. IEEE, 2021.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in neural information processing systems*, 30, 2017.
- Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines, 2019.
- Sangwon Jung, Hongjoon Ahn, Sungmin Cha, and Taesup Moon. Continual learning with node-importance based adaptive group sparse regularization. *Advances in neural information processing systems*, 33:3647–3658, 2020.

- Diederik P Kingma and Max Welling. Auto-encoding variational {Bayes}. In *Int. Conf. on Learning Representations*.
- James Kirkpatrick, Razvan Pascanu, Neil 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. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, mar 2017. doi: 10.1073/pnas.1611835114. URL <https://doi.org/10.1073%2Fpnas.1611835114>.
- Dharshan Kumaran, Demis Hassabis, and James L. McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20(7):512–534, 2016. ISSN 1364-6613. doi: <https://doi.org/10.1016/j.tics.2016.05.004>. URL <https://www.sciencedirect.com/science/article/pii/S1364661316300432>.
- Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 7765–7773, 2018.
- Sergi Masip, Pau Rodriguez, Tinne Tuytelaars, and Gido M van de Ven. Continual learning of diffusion models with generative distillation. *arXiv preprint arXiv:2311.14028*, 2023.
- James McClelland, Bruce McNaughton, and Randall O’Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102:419–57, 08 1995. doi: 10.1037/0033-295X.102.3.419.
- James L. McClelland, Bruce L. McNaughton, and Randall C. O’Reilly. Complementary learning systems within the hippocampus: A neural network modeling approach to understanding episodic memory consolidation. *Psychological Review*, 102(3):419–457, 1995.
- Michael McCloskey and Neal J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. volume 24 of *Psychology of Learning and Motivation*, pp. 109–165. Academic Press, 1989. doi: [https://doi.org/10.1016/S0079-7421\(08\)60536-8](https://doi.org/10.1016/S0079-7421(08)60536-8). URL <https://www.sciencedirect.com/science/article/pii/S0079742108605368>.
- Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- Amy N Myers and Herbert S Wilf. Some new aspects of the coupon collector’s problem. *SIAM review*, 48(3):549–565, 2006.
- Donald J. Newman. The double dixie cup problem. *The American Mathematical Monthly*, 67(1):58–61, 1960. ISSN 00029890, 19300972. URL <http://www.jstor.org/stable/2308930>.
- Andrew Y. Ng and Stuart J. Russell. Inverse reinforcement learning. In *Proceedings of the 17th International Conference on Machine Learning (ICML-2000)*, pp. 663–670, 2000.
- Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. In *International Conference on Learning Representations*, 2018.
- Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.
- Randall C. O’Reilly and Kenneth A. Norman. Hippocampal and neocortical contributions to memory: Advances in the complementary learning systems framework. *Trends in Cognitive Sciences*, 6(12):505–510, December 2002. ISSN 1364-6613. doi: 10.1016/S1364-6613(02)02005-3.
- Mark Ring. *Continual Learning in Reinforcement Environments*. PhD thesis, University of Texas at Austin, 1994. URL <https://www.cs.utexas.edu/~ring/Ring-dissertation.pdf>.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pp. 234–241. Springer, 2015.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016.

- 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 *International conference on machine learning*, pp. 4528–4537. PMLR, 2018.
- Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30, 2017.
- James Seale Smith, Yen-Chang Hsu, Lingyu Zhang, Ting Hua, Zsolt Kira, Yilin Shen, and Hongxia Jin. Continual diffusion: Continual customization of text-to-image diffusion with c-lora. *arXiv preprint arXiv:2304.06027*, 2023.
- Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. *CoRR*, abs/1503.03585, 2015. URL <http://arxiv.org/abs/1503.03585>.
- Gido M. van de Ven and Andreas S. Tolias. Three scenarios for continual learning, 2019.
- Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- Maciej Wołczyk, Michał Zajac, Razvan Pascanu, Łukasz Kuciński, and Piotr Miłoś. Continual world: A robotic benchmark for continual reinforcement learning. *Advances in Neural Information Processing Systems*, 34:28496–28510, 2021.
- Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In *6th International Conference on Learning Representations, ICLR 2018*. International Conference on Learning Representations, ICLR, 2018.
- 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*, pp. 1094–1100. PMLR, 2020.
- Michał Zajac, Kamil Deja, Anna Kuzina, Jakub M. Tomczak, Tomasz Trzcinski, Florian Shkurti, and Piotr Miłoś. Exploring continual learning of diffusion models. *ArXiv*, abs/2303.15342, 2023.

APPENDIX

A HYPERPARAMETERS

A.1 FINETUNE

Hyperparameter	Value	Brief Description
batch size	32	number of samples in each training iteration
epochs	250	number of times the entire dataset is passed through per task
learning rate	10^{-4}	learning rate for gradient descent
optimization algorithm	Adam	optimization algorithm used
β_1	0.9	exponential decay rate for first moment estimates in Adam
β_2	0.999	exponential decay rate for second moment estimates in Adam
epsilon	10^{-8}	small constant for numerical stability
weight decay	0	weight regularization

A.2 MULTITASK

Hyperparameter	Value	Brief Description
batch size	32	number of samples in each training iteration
epochs	500	number of times the entire dataset of all tasks is passed through
learning rate	10^{-4}	learning rate for gradient descent
optimization algorithm	Adam	optimization algorithm used
β_1	0.9	exponential decay rate for first moment estimates in Adam
β_2	0.999	exponential decay rate for second moment estimates in Adam
epsilon	10^{-8}	small constant for numerical stability
weight decay	0	weight regularization

A.3 oEWC

Hyperparameter	Value	Brief Description
batch size	32	number of samples in each training iteration
epochs	250	number of times the entire dataset of all tasks is passed through
learning rate	10^{-4}	learning rate for gradient descent
Fisher multiplier	10^2	the Fisher is scaled by this number to form the EWC penalty
optimization algorithm	Adam	optimization algorithm used
β_1	0.9	exponential decay rate for first moment estimates in Adam
β_2	0.999	exponential decay rate for second moment estimates in Adam
epsilon	10^{-8}	small constant for numerical stability
weight decay	0	weight regularization

The Fisher multiplier hyperparameter was tuned with the values: $10^{-2}, 10^{-1}, 10^0, 10^1, 10^2, 10^3, 10^4, 10^5, 10^6$. We selected the value 10^2 based on the success rate metric given by Equation 2.

A.4 PACKNET

Hyperparameter	Value	Brief Description
batch size	32	number of samples in each training iteration
epochs	250	number of times the entire dataset of all tasks is passed through
retrain epochs	125	number of training epochs after pruning
learning rate	10^{-4}	learning rate for gradient descent
prune percent	0.75	percent of free parameters pruned for future tasks
optimization algorithm	Adam	optimization algorithm used
β_1	0.9	exponential decay rate for first moment estimates in Adam
β_2	0.999	exponential decay rate for second moment estimates in Adam
epsilon	10^{-8}	small constant for numerical stability
weight decay	0	weight regularization

The retrain epochs and prune percent hyperparameters were chosen following the approach in the original PackNet paper. After training the first task, bias layers are frozen.

A.5 DGR

Hyperparameter	Value	Brief Description
batch size	32	number of samples in each training iteration
epochs	250	number of times the entire dataset of all tasks is passed through
learning rate	10^{-4}	learning rate for gradient descent
diffusion training steps	10^4	number of training steps for the diffusion model per task
diffusion warmup steps	$5 * 10^4$	number of extra training steps for the diffusion model on the first task
diffusion timesteps	10^3	number of timesteps in the diffusion process
replay ratio	0.9	percentage of training examples that are generated
optimization algorithm	Adam	optimization algorithm used
β_1	0.9	exponential decay rate for first moment estimates in Adam
β_2	0.999	exponential decay rate for second moment estimates in Adam
epsilon	10^{-8}	small constant for numerical stability
weight decay	0	weight regularization

A.6 CRIL

Hyperparameter	Value	Brief Description
batch size	32	number of samples in each training iteration
epochs	300	number of times the entire dataset of all tasks is passed through
learning rate	10^{-4}	learning rate for gradient descent
diffusion training steps	10^4	number of training steps for the diffusion model per task
diffusion warmup steps	$5 * 10^4$	number of extra training steps for the diffusion model on the first task
diffusion timesteps	10^3	number of timesteps in the diffusion process
replay ratio	0.9	percentage of training examples that are generated
optimization algorithm	Adam	optimization algorithm used
β_1	0.9	exponential decay rate for first moment estimates in Adam
β_2	0.999	exponential decay rate for second moment estimates in Adam
epsilon	10^{-8}	small constant for numerical stability
weight decay	0	weight regularization

A.7 T-DGR

Hyperparameter	Value	Brief Description
batch size	32	number of samples in each training iteration
epochs	300	number of times the entire dataset of all tasks is passed through
learning rate	10^{-4}	learning rate for gradient descent
diffusion training steps	10^4	number of training steps for the diffusion model per task
diffusion warmup steps	$5 * 10^4$	number of extra training steps for the diffusion model on the first task
diffusion timesteps	10^3	number of timesteps in the diffusion process
replay ratio	0.9	percentage of training examples that are generated
optimization algorithm	Adam	optimization algorithm used
β_1	0.9	exponential decay rate for first moment estimates in Adam
β_2	0.999	exponential decay rate for second moment estimates in Adam
epsilon	10^{-8}	small constant for numerical stability
weight decay	0	weight regularization

B MODEL ARCHITECTURE

B.1 MULTI-LAYER PERCEPTRON

Layer (type)	Output Shape	Param #
Linear-1	[32, 512]	25,600
Linear-2	[32, 512]	262,656
Linear-3	[32, 512]	262,656
Linear-4	[32, 512]	262,656
Linear-5	[32, 4]	2,052
Total params:		815,620
Trainable params:		815,620
Non-trainable params:		0

Table 2: Multi-layer perceptron architecture of the learner shared by all methods, featuring five linear layers with ReLU activations applied after each layer except the final one.

To condition the MLP learner on the task, the input vector is concatenated with a one-hot vector representing the task before being passed into the MLP.

B.2 DIFFUSION U-NET

The U-net consists of 4 downsampling and 4 upsampling layers. To condition the generator U-net on the trajectory timestep, we use a sinusoidal positional encoder from Transformers to create an sinusoidal positional embedding that gets passed through 3 linear layers and then added to the convolutional layer output at each level of the U-net.

B.3 GENERATIVE ADVERSARIAL NETWORK

The generator is implemented with an MLP consisting of 7 linear layers with batch normalization and Leaky ReLU activations applied after each layer except the final one.

The discriminator is implemented with an MLP consisting of 3 linear layers with Leaky ReLU activations applied after each layer except the final one.

Sinusoidal positional encodings are fed as input to a 2-layer MLP to condition both the generator and discriminator on the trajectory timestep for t-DGR.

C EXPERIMENT DETAILS

We utilized the following random seeds for the experiments: 1, 2, 3, 4, 5. All experiments were conducted on Nvidia A100 GPUs with 80 GB of memory. The computational node consisted of an Intel Xeon Gold 6342 2.80GHz CPU

with 500 GB of memory. For our longest benchmark, CW20, the runtimes were as follows: DGR and t-DGR took 3 days, CRIL took 16 hours, finetune and oEWC took 6 hours, and PackNet took 8 hours.