

ProtoCRL: Prototype-based Network for Continual Reinforcement Learning

Anonymous authors

Paper under double-blind review

Keywords: Continual reinforcement learning, experience replay, event tables, prototype-based architecture, Gaussian mixture model, variational inference

Summary

The purpose of continual reinforcement learning is to train an agent on a sequence of tasks such that it learns the ones that appear later in the sequence while retaining the ability to perform the tasks that appeared earlier. Experience replay is a popular method used to make the agent remember previous tasks, but its effectiveness strongly relies on the selection of experiences to store. [Kompella et al. \(2023\)](#) proposed organizing the experience replay buffer into partitions, each storing transitions leading to a rare but crucial event, such that these key experiences get revisited more often during training. However, the method is sensitive to the manual selection of event states. To address this issue, we introduce ProtoCRL, a prototype-based architecture leveraging a variational Gaussian mixture model to automatically discover effective event states and build the associated partitions in the experience replay buffer. The proposed approach is tested on a sequence of MiniGrid environments, demonstrating the agent’s ability to adapt and learn new skills incrementally.

Contribution(s)

1. This paper introduces ProtoCRL, a prototype-based architecture for continual reinforcement learning. ProtoCRL features a variational Gaussian mixture model to automatically identify effective event states and build the associated event tables, i.e., partitions within the experience replay buffer (ERB) storing transitions that lead to a particular event state.
Context: Experience replay is a common strategy used in continual reinforcement learning ([Liotet et al., 2022](#); [Luo et al., 2023](#)). [Kompella et al. \(2023\)](#) showed that partitioning the ERB into event tables increases sample efficiency and improves the agent’s generalization performance. However, the method is sensitive to the manual selection of event states. ProtoCRL automatizes the construction of the ERB, making event tables suitable to applications in which the identification of event states is nontrivial.
2. The learned Gaussian mixture components practically serve as prototypical representations of an event state. By inspecting the assignments of the input experiences to the Gaussian mixture components, we show that ProtoCRL identifies meaningful event states that the agent needs to visit more often to remember previously learned tasks.
Context: In the literature, prototypes have been used to either explain pre-trained black-box agents ([Borzillo et al., 2023](#)) or to improve the agents generalization performance of agents trained on single tasks ([Liu et al., 2023](#)). In this work, we leverage the learned prototypical representations to both guide experience replay and gain insights into what information is useful for the agents to maintain the ability to perform multiple tasks learned in sequence.
3. We show that ProtoCRL achieves comparable performance to manually defined event tables and even higher performance when reducing the ERB capacity.
Context: We test ProtoCRL on a sequence of three MiniGrid environments ([Chevalier-Boisvert et al., 2018](#)), comparing its performance in terms of average return and forgetting to manually defined event tables and to ContinualDreamer ([Kessler et al., 2023](#)).

ProtoCRL: Prototype-based Network for Continual Reinforcement Learning

Anonymous authors

Paper under double-blind review

Abstract

1 The purpose of continual reinforcement learning is to train an agent on a sequence of
 2 tasks such that it learns the ones that appear later in the sequence while retaining the
 3 ability to perform the tasks that appeared earlier. Experience replay is a popular method
 4 used to make the agent remember previous tasks, but its effectiveness strongly relies on
 5 the selection of experiences to store. [Kompella et al. \(2023\)](#) proposed organizing the
 6 experience replay buffer into partitions, each storing transitions leading to a rare but
 7 crucial event, such that these key experiences get revisited more often during training.
 8 However, the method is sensitive to the manual selection of event states. To address this
 9 issue, we introduce ProtoCRL, a prototype-based architecture leveraging a variational
 10 Gaussian mixture model to automatically discover effective event states and build the
 11 associated partitions in the experience replay buffer. The proposed approach is tested
 12 on a sequence of MiniGrid environments, demonstrating the agent’s ability to adapt and
 13 learn new skills incrementally.

1 Introduction

15 Recent advances in deep Reinforcement Learning (RL) have achieved super-human performance in
 16 several applications ([Silver et al., 2016](#); [Wurman et al., 2022](#)). However, unlike humans, RL agents
 17 lack the ability to continuously and incrementally learn new skills, which is particularly crucial in
 18 open-world games featuring diverse challenges. In these settings, the major obstacle to continual
 19 learning is catastrophic forgetting ([McCloskey & Cohen, 1989](#); [French, 1999](#)), a phenomenon con-
 20 sisting of new knowledge overwriting previously learned information. Continual RL (CRL) aims to
 21 address this issue by proposing methods to mitigate forgetting and improve the agent’s performance
 22 in sequential tasks ([Khetarpal et al., 2022](#)).

23 Experience Replay (ER) is a common strategy used in both RL and CRL, albeit with slightly differ-
 24 ent objectives. In both contexts, ER involves storing the agent’s experiences in an Experience Replay
 25 Buffer (ERB) and randomly sampling from it to update the agent’s policy. In standard single-task
 26 RL settings, ER thus leads to better generalization and higher stability, as it ensures that learning
 27 relies on a more diverse set of experiences and breaks the correlation between consecutive experi-
 28 ences. In CRL, instead, the ERB also stores data from previous tasks, which is periodically replayed
 29 to prevent catastrophic forgetting ([Rolnick et al., 2019b](#)).

30 In environments where crucial events occur rarely, [Kompella et al. \(2023\)](#) propose to partition the
 31 ERB into event tables, each storing transitions leading to a particular event, i.e. a rarely visited
 32 state that is essential for task completion. Their approach, called Stratified Sampling from Event
 33 Tables (SSET), demonstrates increased robustness to perturbations compared to other ER strategies,
 34 showcasing its effectiveness in mitigating catastrophic forgetting. However, manual selection of
 35 meaningful event states might not always be straightforward, especially in complex games in which
 36 the environment has many diverse features.

To overcome this limitation and to ensure a more flexible approach for continual reinforcement learning, we introduce ProtoCRL, the first prototype-based architecture for CRL. By relying on a Variational Gaussian Mixture Model (VGMM), ProtoCRL automatically identifies event states, assigning each experience to one of the VGMM components, thus eliminating the need for manual event definitions. Moreover, given that the VGMM can learn to use fewer clusters than its maximum capacity, we just need to define the maximum number of components and it will autonomously learn to use additional clusters, and thus additional event tables, for new tasks as they arise. When evaluated on three sequential MiniGrid tasks, ProtoCRL not only reaches similar performance to baselines that rely on predefined event tables, but also surpasses them under memory constraints. Moreover, inspecting the automatically built event tables offers better insight into the agent’s decision-making. Finally, although we use ProtoCRL as online and target network in DDQN (Hasselt et al., 2016), its modular design makes it easily adaptable to other algorithms.

Our main contributions are:

- Proposing the first prototype-based architecture featuring a VGMM for automatic event discovery. ProtoCRL additionally makes event tables more flexible and thus suited to CRL settings, as the VGMM can dynamically leverage fewer clusters than the maximum number given, leaving extra capacity to accommodate additional event tables as new tasks appear.
- Showing that ProtoCRL identifies meaningful event states by inspecting the assignments of observations to the VGMM’s components, thus gaining insights into the agent’s decision process.
- Providing evidence of ProtoCRL comparable performance to baselines with manually defined event tables and superior performance in the presence of memory constraints on a sequence of three MiniGrid tasks.

2 Preliminaries

Reinforcement Learning In this work, we consider an agent acting in a Partially Observable Markov Decision Process (POMDP) (Kaelbling et al., 1998) $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, R, \Omega, O, \gamma \rangle$, with state space \mathcal{S} , action space \mathcal{A} , transition function $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow Pr[\mathcal{S}]$, reward function $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, observation function $O : \mathcal{S} \times \mathcal{A} \rightarrow Pr[\Omega]$, where Ω represents a finite set of observations that the agent can experience, and discount factor $\gamma \in [0, 1)$. The partial observability comes from the fact that the agent does not have access to the environment states $s \in \mathcal{S}$, but only to observations $o \in \Omega$. At time step t , the agent uses its current policy $\pi_t : \Omega \rightarrow Pr[\mathcal{A}]$ to choose an action and observes the immediate reward $r_t \sim R(s_t, a_t)$ and the next observation $o' \sim O(s_t, a_t)$. The episode ends if a termination condition is verified or if a horizon of H is reached. In the remaining part of this introduction, we assume for convenience that the agent has access to states s_t . In the partial observability case, all mathematical definitions would be the same but with o in place of s . We return to partial observability in Section A.1 of the *Supplementary materials*, where we describe the exact form and content of these observations.

The value function of a policy is represented by its long-term discounted return: $Q^\pi(s, a) = R(s, a) + \gamma E_{s' \sim \mathcal{P}(s, a)}[V^\pi(s')]$, where $V^\pi(s) = Q^\pi(s, \pi(s))$. The aim of the agent is to find an optimal policy π^* and the corresponding $Q^*(s, a)$ that maximizes the expected discounted return. Model-free off-policy methods learn $Q^*(s, a)$ directly from data by incrementally updating the value function based on the temporal difference (TD)-error $\delta = r(s, a) + \gamma V_k(s') - Q_k(s, a)$ for $V_k(s') = \max_{a'} Q_k(s', a')$. DDQN (Hasselt et al., 2016) is one such approach, in which the value function is represented as a deep neural network with parameters ϕ , that are updated along the gradient of the TD-error. With respect to DQN (Mnih et al., 2015), DDQN uses two different value networks: an online network for current Q-values and a target network for computing $V_k(s')$ that is updated with lower frequency.

Experience Replay with Event Tables In off-policy RL, experiences $\langle s, a, r, s' \rangle$ are typically stored in an ERB and then sampled in batches to perform gradient updates to the value networks

parameters. [Kompella et al. \(2023\)](#) propose partitioning the ERB into n event tables \mathcal{B}^{ν_i} and a default table \mathcal{B}^0 . The default table stores all transitions, while each event table is associated with an event specification $\nu = \langle \omega, \tau \rangle$, where $\omega : \mathcal{S} \rightarrow \{0, 1\}$ is a boolean event condition and τ is a history length. Typical good candidates for event states are important states that occur rarely, such as goal states or bottleneck states, and their associated event conditions need to be specified by domain experts. All tables are handled in a FIFO manner and insertion and sampling operations are controlled by parameters η, κ, d_i , with $i \in [0, n]$, representing sampling probabilities, capacity sizes, and minimum data requirements for each table (including the default table, corresponding to $i = 0$), respectively.

Continual Reinforcement Learning CRL extends the standard RL paradigm to scenarios where an agent is faced with a sequence of tasks or changing environments. Each new task is represented as a new POMDP $\mathcal{M}_i = \langle \mathcal{S}_i, \mathcal{A}_i, \mathcal{P}_i, R_i, \Omega_i, O_i, \gamma_i \rangle$, with $i = 1, \dots, T$. During the i -th task training, the agent collects only experiences from that task.

3 Related Work

Continual Reinforcement Learning In recent years, interest in CRL has grown considerably ([Khetarpal et al., 2022](#)). However, the study of catastrophic forgetting and the more general stability-plasticity dilemma has a much longer history ([McCloskey & Cohen, 1989](#); [French, 1999](#); [Mermillod et al., 2013](#)). While some approaches aim at reducing forgetting by applying some form of weight or function regularization ([Kirkpatrick et al., 2016](#); [Kessler et al., 2022](#)), architecture-based strategies feature task-specific sub-networks or dynamic networks to prevent task interference ([Rusu et al., 2016](#); [Mallya & Lazebnik, 2017](#)). Another popular solution is experience replay (ER), which consists of storing and replaying experiences from old tasks ([Rolnick et al., 2019a](#); [Li et al., 2021](#); [Liotet et al., 2022](#); [Luo et al., 2023](#)). For instance, CLEAR ([Rolnick et al., 2019a](#)) balances plasticity and stability by alternating off-policy learning from the ERB and on-policy learning from new experiences, using behavioural cloning between the current policy and its past version to increase stability. The way in which we populate and sample from the ERB significantly affects the final performance, as demonstrated by [Kessler et al. \(2023\)](#). In particular, ER is not particularly suited in environments featuring crucial event states that occur very rarely during each episode. These states might get sampled too infrequently from the ERB, leading to the agent being unable to efficiently and effectively learn the task at hand. [Kompella et al. \(2023\)](#) address this issue by dividing the ERB into event tables and ensuring they get sampled frequently enough. However, the effectiveness of event tables strongly relies on the definition of event conditions by domain experts, and a wrong choice of events might lead to suboptimal performance. The purpose of ProtoCRL is to avoid this problem by automatizing the construction of event tables.

Prototypes in Reinforcement Learning The term *prototype* has divergent meanings across different research sub-fields. In representation learning, a prototype is a latent vector encoding features relevant for achieving a certain outcome, i.e., a prototypical embedding in the latent space. In explainable artificial intelligence, a prototype is typically an input instance or a part of it (e.g., an image patch) that is representative of a specific class.

While early prototype-based networks focused on supervised image classification ([Chen et al., 2019](#); [Rymarczyk et al., 2022](#)), recent works in RL adapt these principles to interpret black-box image-based RL agents. [Ragodos et al. \(2022\)](#) exploit the black-box model’s demonstrations to train a self-interpretable agent by imitation learning, while [Borzillo et al. \(2023\)](#) and [Kenny et al. \(2023\)](#) use the pre-trained black-box agent as an encoder. In representation learning, instead, self-supervised learning and task-agnostic pre-training allow learning prototypical embeddings to enhance the generalization capabilities of RL agents ([Yarats et al., 2021](#); [Liu et al., 2023](#); [Mazouze et al., 2022](#)). For instance, DreamerPro ([Deng et al., 2022](#)) clusters observations into trainable prototypes and predicts the cluster assignment from both the world model’s state and an augmented view of the

133 observations, while ProtoCAD (Wang et al., 2024) utilizes prototypical representations to extract
 134 contextual information from varying dynamic environments.

135 To our knowledge, no existing work applies a prototype-based network to CRL. We address this
 136 gap via an architecture featuring a VGMM (Corduneanu & Bishop, 2001; Nasios & Bors, 2006)
 137 for clustering latent representations and identifying prototypical embeddings of event states. Unlike
 138 prior methods, our architecture does not require a pre-trained black-box agent or self-supervised
 139 pre-training.

140 4 ProtoCRL

141 In this section, we introduce ProtoCRL, our prototype-based architecture that addresses the chal-
 142 lenges of CRL by integrating a prototype-based latent representation with an event-based experience
 143 replay mechanism. By learning to cluster latent features via a VGMM, ProtoCRL automatically
 144 identifies prototypical event states that are critical for learning. These prototypes are then used to
 145 form event tables in the replay buffer, ensuring that rare but informative transitions are frequently
 146 revisited during training.

147 4.1 Network Architecture

148 The network, illustrated in Figure 1, consists of three interdependent components: an encoder f_e , a
 149 VGMM module, and an output layer f_o .

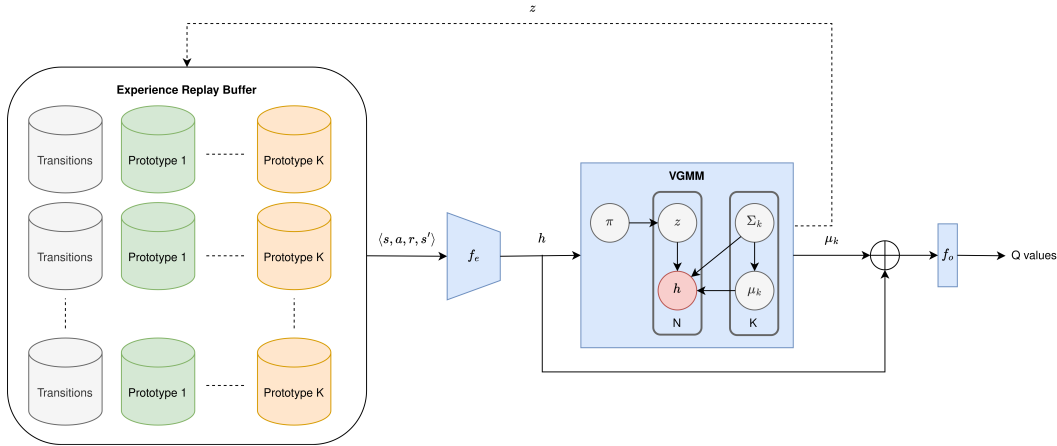


Figure 1: **ProtoCRL**. A transition $\langle s, a, r, s' \rangle$ is sampled from the ERB and state s is passed to the encoder f_e , which produces latent representation h . h is the input to a VGMM with K components. The VGMM assigns the input observation to one of its components through $z \in \mathcal{Z}^K$, which is a one-hot assignment vector. The VGMM is parametrized by mixing coefficients $\pi = (\pi_1, \dots, \pi_K)$, means $\mu = (\mu_1, \dots, \mu_K)$, and covariance matrices $\Sigma = (\Sigma_1, \dots, \Sigma_K)$. The latent state representation h and the mean of the component it is assigned to μ_k are concatenated and passed to an output dense layer f_o , which predicts Q values. When collecting the experiences to store in the ERB, each experience is stored in the prototypical event table associated with the VGMM component the observation is assigned to.

150 **Encoder** The encoder f_e transforms high-dimensional raw observations into a compact latent rep-
 151 resentation. This component can be implemented differently, depending on the type of input ob-
 152 servations. The learned latent space is crucial for capturing the relevant features needed for both
 153 control and for subsequent clustering. In our experiments, we implement the encoder as a simple
 154 dense neural network with three hidden layers and LeakyReLU nonlinearities.

Variational Gaussian Mixture Model (VGMM) The VGMM processes the latent representations produced by the encoder and clusters them into a set of prototypes. A *Gaussian Mixture Model* (GMM) is a widely-used probabilistic model for clustering and density estimation. The GMM assumes that the observed data is generated from a mixture of several Gaussian distributions, each representing a different cluster or component in the data. In this work, we deal with multi-dimensional data, so we use a *multivariate GMM* which models the distribution of a set of N observed data points $\mathcal{X} = x_1, \dots, x_N$, with $x_n \in \mathbb{R}_d$, as a weighted sum of K multivariate Gaussian components. The probability density function for the multivariate GMM is given by:

$$P(x_n|\theta) = \sum_{k=1}^K \pi_k \mathcal{N}(x|\mu_k, \Sigma_k)$$

where π_k are the mixing weights, $\mathcal{N}(x_n|\mu_k, \Sigma_k)$ is a multivariate Gaussian distribution with mean μ_k and covariance matrix Σ_k , and $\theta = \{\pi_k, \mu_k, \Sigma_k\}_{k=1}^K$ represents the set of all parameters of the model. The mixing weights are such that $0 \leq \pi_k \leq 1$ and $\sum_{k=1}^K \pi_k = 1$, as they represent the probability that a data point belongs to each mixture component. Given our GMM, the end goal is to perform inference, i.e., to estimate the parameters of our mixture components given some observations:

$$p(\theta|\mathcal{X}) = \frac{p(\mathcal{X}|\theta)p(\theta)}{p(\mathcal{X})}$$

where $p(\mathcal{X}|\theta)$ is the likelihood of the data given the parameters, and $p(\theta)$ is the prior distribution. However, the computation of the denominator, called evidence, is generally intractable. The main idea of variational inference is to find a simple distribution $q(\theta, \phi)$ from a family of distributions \mathcal{Q} , called variational distributions, that is a good approximation of the true posterior. To this aim, we look for the distribution $q^*(\theta, \phi)$ that minimizes the Kullback-Leibler (KL) divergence between the true and approximate posteriors, which is equivalent to maximizing the evidence lower bound (ELBO):

$$q^*(\theta, \phi) = \arg \min_{q(\theta, \phi) \in \mathcal{Q}} D_{KL} [q(\theta, \phi) || p(\theta, \phi|\mathcal{X})]$$

155 To this aim, we need to place prior distributions over the model parameters θ . In this work, we use
 156 a Dirichlet prior for the mixing coefficients π_k , a Gaussian prior for the means μ_k , and a Gamma
 157 prior for the magnitude of the precision matrix, as reported in the literature (Lu, 2021).

To make this optimization scalable, especially for large datasets, we use Stochastic Variational Inference (SVI) (Hoffman et al., 2013). At each iteration, the gradient of the ELBO with respect to the variational parameters ϕ is estimated using the batch of latent representations produced by the encoder. The gradient is then used to update ϕ using a learning rate η_t , following the update rule:

$$\phi_{t+1} = \phi_t + \eta_t \nabla_{\phi} ELBO(\phi_t, \mathcal{X}_{mini-batch})$$

158 By iterating over mini-batches, the variational parameters ϕ converge to values that approximate
 159 the true posterior distribution of the model parameters. The main advantage of using variational
 160 inference is that after defining the maximum number of components in the VGMM, a smaller number
 161 of components might be used based on the input data. This is particularly useful in our CRL setting,
 162 in which an increasing number of components are needed as the number of seen tasks grows.

163 **Output Layer** The output layer maps the prototype-augmented features, obtained by concatenat-
 164 ing the encoder’s output with the mean of the GMM component it is assigned to, to action-value
 165 estimates using a shared network head. In our continual learning setting, the use of a shared head
 166 across tasks promotes knowledge transfer while reducing the overall model complexity. However,
 167 this choice may also introduce interference if tasks are highly divergent, which is mitigated by our
 168 selective replay strategy.

169 4.2 Prototypical Event Tables

To enhance learning and avoid catastrophic forgetting, ProtoCRL employs event tables (Kompella et al., 2023). Specifically, for each task $i \in [1, T]$, we define a default table \mathcal{B}_i^0 and prototypical event tables $\mathcal{B}_i^{\nu_k}$, with $k \in [1, K]$. Let $h = f_e(s)$ be the latent representation of state s obtained from the encoder. For a VGMM with K components, if we denote the posterior probability that h is generated by component k as $p(k | h)$, the event condition for table k can be defined as the indicator function

$$\omega_k(s) = \mathbf{1} \left\{ k = \arg \max_{j \in \{1, \dots, K\}} p(j | f_e(s)) \right\}$$

170 which evaluates to 1 if and only if the latent representation $h = f(s)$ is most likely generated by
 171 the k -th component, and 0 otherwise. By over-sampling transitions from these event tables during
 172 training, the network more frequently reinforces critical experiences, which is particularly valuable
 173 in the context of continual learning. One key aspect in the construction of manual event tables
 174 is the definition of the sampling probabilities η_i , which is crucial for guaranteeing their effective-
 175 ness. For simplicity, we fix the same sampling probability for each event table throughout training,
 176 leaving more sophisticated tuning strategies for future work. However, results show that ProtoCRL
 177 effectively manages to distribute states among tables to favor more frequent sampling of important
 178 states.

179 4.3 Training Procedure

VGMM Warm Start To ensure initial reliable clustering in the latent space, the VGMM is pre-trained (or warm-started) on data collected by performing a random policy on the first task in the sequence for 200 steps. Although the agent has yet to see future tasks, this warm start builds a reasonable basis for prototype extraction. As training proceeds through subsequent tasks, the clustering continues to be updated in an online fashion to reflect data from the new tasks.. The training objective used during this warm start procedure is made up of several components. The first one is the TD loss, which ensures that the learned representations are meaningful for the task at hand. The second term is the ELBO, that we need to maximize:

$$\mathcal{L}_{ELBO} = \alpha_{logp} \log p(\mathcal{X}) - \alpha_{KL} D_{KL} [q(\theta, \lambda) || p(\theta, \lambda | \mathcal{X})]$$

The third component is an entropy term that encourages the use of diverse clusters. Let $r_i \in \mathbb{R}^K$ denote the cluster probability vector for the i -th sample, obtained by applying a softmax over the VGMM log-probabilities. We define the average probability of cluster k over a batch of N samples as:

$$m_k = \frac{1}{N} \sum_{i=1}^N r_{i,k}, \quad \text{for } k = 1, \dots, K$$

Then, we define the entropy loss as:

$$\mathcal{L}_E = - \sum_{k=1}^K m_k \log(m_k + \epsilon),$$

where ϵ is a small constant for numerical stability. Finally, we use Hoyer’s sparsity loss (Hoyer, 2004) to encourage exactly one VGMM component to have high probability for each input observation. The Hoyer sparsity measure for the i -th input sample takes the form:

$$H(r_i) = \frac{\sqrt{K} \|r_i\|_1}{\|r_i\|_2 + \epsilon} - 1$$

where $\|r_i\|_1 = \sum_{k=1}^K |r_{i,k}|$ and $\|r_i\|_2 = \sqrt{\sum_{k=1}^K r_{i,k}^2}$. A perfectly one-hot vector achieves $H(r_i) = 1$, while a uniform distribution yields a lower value. The overall Hoyer sparsity loss over

a batch of N samples is computed as the mean:

$$\mathcal{L}_H = \frac{1}{N} \sum_{i=1}^N H(r_i).$$

Minimizing this loss encourages each probability vector r_i to become more sparse. The overall training objective is:

$$\mathcal{L}_{ProtoCRL} = \mathcal{L}_{TD} - \lambda(\mathcal{L}_{ELBO} - \beta_E \mathcal{L}_E + \beta_H \mathcal{L}_H)$$

180 with λ, β_E and β_H being regularization parameters.

181 **Alternating Training Objectives** After the warm start procedure, we alternatively optimize two
 182 objectives: the TD loss, to sharpen the agent’s policy, and the joint objective used during warm
 183 start, to establish a well-structured latent space and meaningful prototypes that form the basis for
 184 effective event table construction. This alternating training strategy is designed to find a balance
 185 between maintaining a robust and interpretable latent structure and optimizing the agent’s decision-
 186 making capabilities. This design not only automates the discovery of significant event states but also
 187 improves the overall stability and performance of the agent in continual learning environments.

188 5 Experiments

189 We test ProtoCRL on a sequence of 3 MiniGrid (Chevalier-Boisvert et al., 2018) tasks - *DoorKey*,
 190 *LavaCrossing*, and *SimpleCrossing* - using DDQN algorithm (Hasselt et al., 2016). The environ-
 191 ments are episodic and terminate when the agent reaches the goal or when we reach the maximum
 192 episode length (set to 200 at evaluation time). The agent receives a reward of -0.05 for each
 193 environment step, $+1$ for reaching the goal, and -5 for falling into lava. Further details on the
 194 environments, architectures and hyperparameters needed for reproducibility are given in Section A
 195 of the *Supplementary Materials*. We train all agents for 3K epochs (1K epochs for each task) and
 196 evaluate them with a frequency of 20 epochs, averaging results over 5 seeds. For each experiment,
 197 we report average performance and average forgetting CRL metrics (Wołczyk et al., 2021), whose
 198 formal definitions are provided in the *Supplementary materials*. The former measures how well the
 199 agents performs on average on all tasks at the end of the training, while the latter captures the drop
 200 in the agent’s performance on previous tasks after learning the entire task sequence.

201 A central goal of ProtoCRL is to eliminate the need for manually selecting event states while pre-
 202 serving performance in CRL settings. For this reason, we compare it to (1) a baseline using a single
 203 FIFO ERB with no event tables (NoET), (2) a simplified variant of event-based replay in which
 204 we introduce a single event table holding transitions related to the *goal* state, (3) an event-based
 205 strategy using multiple event tables (*goal*, *at_door*, *pickup_key* for *DoorKey*, *goal* and *at_lava*
 206 for *LavaCrossing*, and *goal* and *at_gap* for *SimpleCrossing*), (4) ContinualDreamer (Kessler et al.,
 207 2023), which leverages both experience replay and world models. In particular, we expect Continu-
 208 alDreamer to be a strong baseline, as its learned dynamic model can generate additional transitions
 209 and thus provide more frequent revisits of past tasks.

210 5.1 Performance Comparison

211 In this section, we compare the agents performance when using an ERB with capacity of 2M. Note
 212 that event tables are partitions within the ERB, therefore the total memory usage remains fixed at 2M
 213 transitions, regardless of the number of event tables. Figure 2 shows episodic returns for all tasks
 214 across the whole training. Looking at the plot on the left, focusing on the *DoorKey* task, we notice
 215 that all methods steadily improve over the training epochs, but experience different performance
 216 drops when encountering the second task in the sequence, at epoch 1K. While ProtoCRL, GoalET,
 217 and ET manage to recover from this drop, NoET keeps losing its ability to perform the *DoorKey* task,
 218 reaching the lowest final performance. In these experiments, we also show that ProtoCRL benefits

from storing a larger history of transitions in the event tables, namely $\tau = 500$ compared to $\tau = 50$, which is the value used for GoalET and ET. Indeed, the second lowest final return is obtained by ProtoCRL-2M-h50, which despite not experiencing forgetting, fails to reach a high performance on the DoorKey task. On the other hand, ContinualDreamer is the approach that converges faster and maintains the highest final return, indicating that the integration of world models and experience replay is well suited for CRL settings.

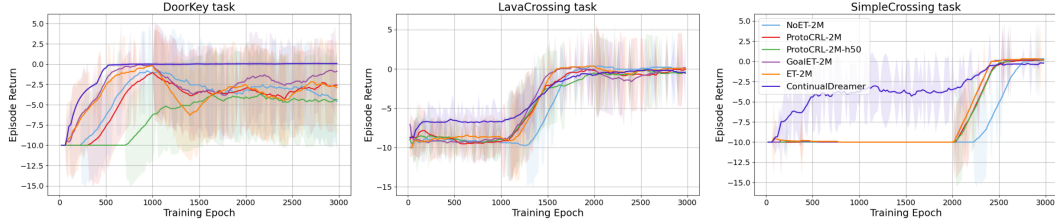


Figure 2: Episodic return on 3 MiniGrid tasks using a 2M memory budget. The task ordering is *DoorKey-LavaCrossing-SimpleCrossing*. Each agent is trained on each task for 1K epochs and evaluated every 20 epochs across the entire training sequence. Shaded areas indicate standard deviations across 5 seeds.

On the subsequent tasks, all methods reach similar performances, slightly above or below 0. Notably, ContinualDreamer is the one obtaining the lowest final return for both LavaCrossing and SimpleCrossing. This results in GoalET having the highest average performance of 0.32 ± 0.02 and the lowest forgetting of -0.01 ± 0.03 , as reported in Table 1. The negative forgetting value indicates that the final performance is higher at the end of the entire training than after training the individual tasks. Despite not being the best-performing method, ProtoCRL reaches a final average performance of -0.42 ± 1.08 that is significantly higher than the NoET and ET baselines, both having average performance below -1 and higher standard deviations across seeds.

Table 1: Average performance and forgetting on three MiniGrid tasks using 2M memory budget.

	Avg. Performance	Avg. Forgetting
NoET	-1.40 ± 2.44	1.72 ± 2.39
GoalET	0.32 ± 0.02	-0.01 ± 0.03
ET	-1.05 ± 1.96	1.30 ± 1.85
ContinualDreamer	0.19 ± 0.07	0.05 ± 0.09
ProtoCRL-h50	-1.07 ± 1.97	0.72 ± 0.97
ProtoCRL	-0.42 ± 1.08	0.73 ± 1.08

5.2 Decreasing Buffer Capacity

Sample efficiency is a crucial requirement in CRL settings. One major disadvantage of experience replay methods is indeed that the size of the memory grows with the number of tasks. It is therefore necessary for CRL methods to perform well even in the presence of memory constraints. For this reason, we test the effectiveness of NoET, GoalET, ET, and ProtoCRL with different buffer capacities.

In Tables 2 and 3, we compare the effectiveness of these methods in terms of average performance and forgetting. In Section C in the *Supplementary materials*, we additionally provide plots showing the episodic returns for the three environments across the entire training. Overall, ET and GoalET are the approaches that suffer from the greatest performance drops (i.e., increased forgetting) when decreasing the memory budget, reaching performances that are much lower than the NoET baseline with 1M and 100K memory budgets. This might be due to manually defined tables oversampling

a narrow subset of transitions, thus limiting the agent’s exposure to broader context. If event definitions or table sampling probabilities are not well tuned, uniform replay (NoET) might maintain a more diverse set of experiences, resulting in better overall performance. On the contrary, ProtoCRL, with its ability to automatically build event tables, manages to achieve high performance with very low memory budgets.

Table 2: Average performance over 3 MiniGrid tasks with different memory budgets. All metrics are an average and standard deviation over 5 seeds. We highlight in bold the best performing approaches.

	100K	500K	1M	2M
NoET	-5.17 ± 4.24	-4.47 ± 4.25	-5.16 ± 3.53	-1.40 ± 2.44
GoalET	-8.61 ± 0.98	-3.80 ± 2.93	-7.24 ± 1.95	0.32 ± 0.02
ET	-9.31 ± 0.97	-4.49 ± 2.57	-7.24 ± 1.95	-1.05 ± 1.96
ProtoCRL	-1.73 ± 2.93	-2.41 ± 0.99	-0.37 ± 0.99	-0.42 ± 1.08

Table 3: Average forgetting over 3 MiniGrid tasks with different memory budgets. All metrics are an average and standard deviation over 5 seeds. We highlight in bold the best performing approaches.

	100K	500K	1M	2M
NoET	4.78 ± 3.48	4.75 ± 4.15	3.43 ± 3.50	0.86 ± 1.21
GoalET	4.08 ± 4.39	2.74 ± 3.88	3.43 ± 4.86	-0.01 ± 0.02
ET	5.47 ± 4.20	3.39 ± 3.44	2.75 ± 3.88	1.33 ± 1.89
ProtoCRL	2.02 ± 1.63	2.01 ± 1.62	0.01 ± 0.01	0.01 ± 1.72

5.3 Prototypes analysis

To shed light on what types of events each prototype, i.e., each VGMM’s mixture component, is representing, we inspect which states or transitions predominantly fall into each cluster. To this aim, we took the agent trained on the entire task sequence and made it perform 1 episode in each environment following the learned policy, and 10 episodes by choosing a random action with probability 70% in order to collect more diverse states, thus collecting 2959 states in total.

In Figure 3, we show the percentage of states assigned to each cluster that are related to specific event states. Specifically, for *goal*, *key*, and *lava* we consider all states in which the agent is in the position of the object or in an adjacent position but oriented towards the object. For *door* and *gap*, instead, we consider only states in which the agent is exactly at the door and gap positions.

Interestingly, around 25% of the states in Cluster 3 represent *key*, *door*, and *goal* states relevant for the DoorKey task. Additionally, the same cluster has an important coverage of the *gap* state, which is similar to the *door* state, together with smaller percentages of *lava* states and *goal* states for the LavaCrossing and SimpleCrossing tasks. Overall, Cluster 3 appears to be an event cluster, as it captures many state transitions that are important to the agent’s task progression. Similarly, around 17% of the states in Cluster 5 seem to be capturing event-related states. Lava-related events, instead, are predominantly collected in Cluster 1, although there was 1 state among the 2959 we collected that was the only one assigned to Cluster 9. Because no cluster stands out with more than 50% of its states representing a particular event, it appears that the VGMM’s components might be capturing more nuanced distinctions (e.g., each cluster is partially capturing key states under different local configurations, such as the agent’s orientation or position near the key) or sub-trajectories leading to overcoming the obstacles or reaching the goal. This might seem the case by looking at Figure 4, which shows an overlay of the states visited by the agent’s policy that are assigned to each cluster. In accordance with our previous analysis, Cluster 3 collects the states corresponding to door opening and the subtrajectory leading to the goal in the DoorKey task. Clusters 5 and 8, instead, seems more focused on subtrajectories that allow overcoming obstacles

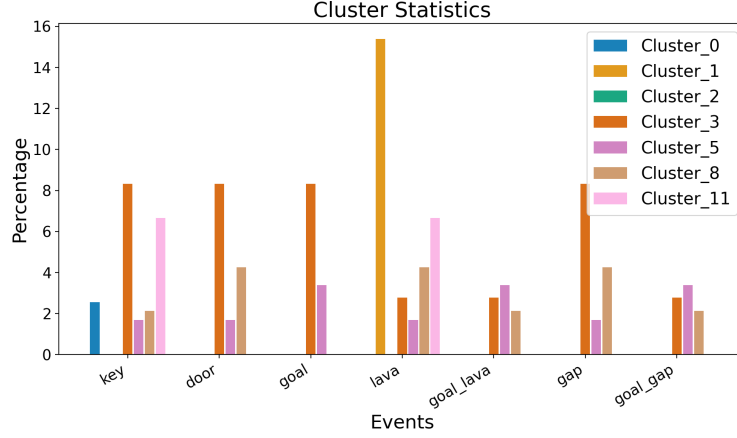


Figure 3: Distribution of environment-defined event types within each discovered prototype cluster. The automatically learned prototypes frequently match to key environment events, demonstrating ProtoCRL’s ability to partition states in an event-centric manner.

276 and reaching the goal in the other two environments. Finally, Cluster 11 collects the state relative
 277 to key collection in the DoorKey task. These results confirm that ProtoCRL effectively identifies
 278 and groups pivotal states for each environment, supporting more efficient experience replay and
 279 removing the need for manually defined event tables.

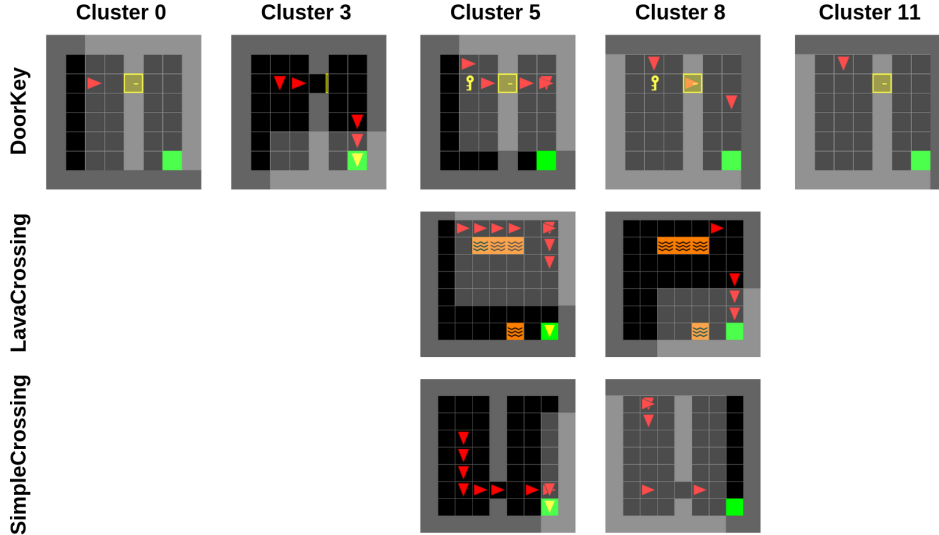


Figure 4: Overlaid states visited by the agent’s policy that are assigned to each cluster.

280 6 Conclusions

281 In this work, we introduce ProtoCRL, a prototype-based network for CRL that automatically discov-
 282 ers event states via a VGMM. By replacing the manual event definitions with an adaptive, data-driven
 283 mechanism for organizing the ERB, ProtoCRL allows the agent to effectively mitigate forgetting.
 284 Our experiments demonstrate that under stringent memory constraints, ProtoCRL outperforms con-
 285 ventional baselines that rely on manually defined event tables or uniform sampling.

While ProtoCRL shows promising sample efficiency and robustness, some limitations remain. The effectiveness of the VGMM-based clustering can be highly sensitive to hyperparameters and may capture nuanced distinctions across subtrajectories rather than isolating single dominant event states. Additionally, the results ProtoCRL achieves are still lower than state-of-the-art CRL approaches integrating experience replay with world models (Kessler et al., 2023). Future work should explore tighter integration between ProtoCRL and world-model-based approaches. Moreover, investigating adaptive hyperparameter schemes that allow learning other sensitive hyperparameters, such as the sampling probabilities for the automatically built event tables, could further enhance the stability and performance of ProtoCRL in more complex settings.

References

- Caterina Borzillo, Alessio Ragno, and Roberto Capobianco. Understanding deep rl agent decisions: a novel interpretable approach with trainable prototypes. In *XAI.it@AI*IA*, 2023. URL <https://api.semanticscholar.org/CorpusID:264491932>.
- Chaofan Chen, Oscar Li, Chaofan Tao, Alina Jade Barnett, Jonathan Su, and Cynthia Rudin. *This looks like that: deep learning for interpretable image recognition*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- Maxime Chevalier-Boisvert, Lucas Willems, and Suman Pal. Minimalistic gridworld environment for openai gym. <https://github.com/maximecb/gym-minigrid>, 2018.
- Adrian Corduneanu and Christopher M Bishop. Variational bayesian model selection for mixture distributions. In *Artificial intelligence and Statistics*, volume 2001, pp. 27–34. Morgan Kaufmann Waltham, MA, 2001.
- Fei Deng, Ingook Jang, and Sungjin Ahn. DreamerPro: Reconstruction-free model-based reinforcement learning with prototypical representations. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 4956–4975. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/deng22a.html>.
- Robert M. French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135, 1999. ISSN 1364-6613. DOI: [https://doi.org/10.1016/S1364-6613\(99\)01294-2](https://doi.org/10.1016/S1364-6613(99)01294-2). URL <https://www.sciencedirect.com/science/article/pii/S1364661399012942>.
- Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, pp. 2094–2100. AAAI Press, 2016.
- Matthew D. Hoffman, David M. Blei, Chong Wang, and John Paisley. Stochastic variational inference. *J. Mach. Learn. Res.*, 14(1):1303–1347, may 2013. ISSN 1532-4435.
- Patrik O. Hoyer. Non-negative matrix factorization with sparseness constraints. *J. Mach. Learn. Res.*, 5:1457–1469, December 2004. ISSN 1532-4435.
- Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artif. Intell.*, 101(1–2):99–134, May 1998. ISSN 0004-3702.
- Eoin M. Kenny, Mycal Tucker, and Julie Shah. Towards interpretable deep reinforcement learning with human-friendly prototypes. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=hWwY_Jq0xsN.

- Samuel Kessler, Jack Parker-Holder, Philip Ball, Stefan Zohren, and Stephen J. Roberts. Same state, different task: Continual reinforcement learning without interference. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(7):7143–7151, Jun. 2022. DOI: 10.1609/aaai.v36i7.20674. URL <https://ojs.aaai.org/index.php/AAAI/article/view/20674>.
- Samuel Kessler, Mateusz Ostaszewski, MichałPaweł Bortkiewicz, Mateusz Źarski, Maciej Wolczyk, Jack Parker-Holder, Stephen J Roberts, Piotr Mi, et al. The effectiveness of world models for continual reinforcement learning. In *Conference on Lifelong Learning Agents*, pp. 184–204. PMLR, 2023.
- Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. Towards continual reinforcement learning: A review and perspectives. *Journal of Artificial Intelligence Research*, 75:1401–1476, 2022.
- 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. *Proceedings of the National Academy of Sciences*, 114:3521 – 3526, 2016. URL <https://api.semanticscholar.org/CorpusID:4704285>.
- Varun Raj Kompella, Thomas Walsh, Samuel Barrett, Peter R. Wurman, and Peter Stone. Event tables for efficient experience replay. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=XejzjAjKjv>.
- Chunmao Li, Yang Li, Yinliang Zhao, Peng Peng, and Xupeng Geng. Sler: Self-generated long-term experience replay for continual reinforcement learning. *Applied Intelligence*, 51(1):185–201, 2021.
- Pierre Liotet, Francesco Vidaich, Alberto Maria Metelli, and Marcello Restelli. Lifelong hyper-policy optimization with multiple importance sampling regularization. In *AAAI Conference on Artificial Intelligence (AAAI)*, pp. 7525–7533. AAAI Press, 2022. URL <https://doi.org/10.1609/aaai.v36i7.20717>.
- Xin Liu, Yaran Chen, Haoran Li, Boyu Li, and Dongbin Zhao. Cross-domain random pre-training with prototypes for reinforcement learning. *arXiv preprint arXiv:2302.05614*, 2023.
- Jun Lu. A survey on bayesian inference for gaussian mixture model. *arXiv preprint arXiv:2108.11753*, 2021.
- Yongle Luo, Yuxin Wang, Kun Dong, Qiang Zhang, Erkang Cheng, Zhiyong Sun, and Bo Song. Relay hindsight experience replay: Self-guided continual reinforcement learning for sequential object manipulation tasks with sparse rewards. *Neurocomputing*, 557:126620, 2023.
- Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7765–7773, 2017. URL <https://api.semanticscholar.org/CorpusID:35249701>.
- Bogdan Mazouze, Ahmed M Ahmed, Patrick MacAlpine, R Devon Hjelm, and Andrey Kolobov. Cross-trajectory representation learning for zero-shot generalization in rl. *International Conference on Learning Representations*, 2022.
- Michael McCloskey and Neal J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation - Advances in Research and Theory*, 24(C):109–165, January 1989. ISSN 0079-7421. DOI: 10.1016/S0079-7421(08)60536-8.
- Martial Mermillod, Aurélia Bugaïska, and Patrick Bonin. The stability-plasticity dilemma: investigating the continuum from catastrophic forgetting to age-limited learning effects. *Frontiers in Psychology*, 4, 2013.

- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Belle-
mare, Alex Graves, Martin A. Riedmiller, Andreas Kirkeby Fidjeland, Georg Ostrovski, Stig
Petersen, Charlie Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran,
Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforce-
ment learning. *Nature*, 518:529–533, 2015. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:205242740)
[CorpusID:205242740](https://api.semanticscholar.org/CorpusID:205242740).
- Nikolaos Nasios and Adrian G Bors. Variational learning for gaussian mixture models. *IEEE*
Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), 36(4):849–862, 2006.
- Ronilo Ragodos, Tong Wang, Qihang Lin, and Xun Zhou. Explaining a reinforcement learning agent
via prototyping. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.),
Advances in Neural Information Processing Systems, 2022. URL [https://openreview.](https://openreview.net/forum?id=nyBJcnhJAoy)
[net/forum?id=nyBJcnhJAoy](https://openreview.net/forum?id=nyBJcnhJAoy).
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience
replay for 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., 2019a. URL [https://proceedings.neurips.cc/paper_](https://proceedings.neurips.cc/paper_files/paper/2019/file/fa7cdfad1a5aaf8370ebeda47a1ff1c3-Paper.pdf)
[files/paper/2019/file/fa7cdfad1a5aaf8370ebeda47a1ff1c3-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/fa7cdfad1a5aaf8370ebeda47a1ff1c3-Paper.pdf).
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience
replay for continual learning. *Advances in neural information processing systems*, 32, 2019b.
- Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirk-
patrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural net-
works. *ArXiv*, abs/1606.04671, 2016. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:15350923)
[CorpusID:15350923](https://api.semanticscholar.org/CorpusID:15350923).
- Dawid Rymarczyk, Łukasz Struski, Michał Górszczak, Koryna Lewandowska, Jacek Tabor, and
Bartosz Zieliński. Interpretable image classification with differentiable prototypes assignment. In
European Conference on Computer Vision, pp. 351–368. Springer, 2022.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche,
Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering
the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Junjie Wang, Qichao Zhang, Yao Mu, Dong Li, Dongbin Zhao, Yuzheng Zhuang, Ping Luo, Bin
Wang, and Jianye Hao. Prototypical context-aware dynamics for generalization in visual control
with model-based reinforcement learning. *IEEE Transactions on Industrial Informatics*, pp. 1–11,
2024. DOI: 10.1109/TII.2024.3396525.
- Maciej Wołczyk, MichałZajac, Razvan Pascanu, Łukasz Kuciński, and Piotr Miłoś. Con-
tinual world: A robotic benchmark for continual reinforcement learning. In M. Ranzato,
A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neu-
ral Information Processing Systems*, volume 34, pp. 28496–28510. Curran Associates, Inc.,
2021. URL [https://proceedings.neurips.cc/paper_files/paper/2021/](https://proceedings.neurips.cc/paper_files/paper/2021/file/ef8446f35513a8d6aa2308357a268a7e-Paper.pdf)
[file/ef8446f35513a8d6aa2308357a268a7e-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/ef8446f35513a8d6aa2308357a268a7e-Paper.pdf).
- Peter R Wurman, Samuel Barrett, Kenta Kawamoto, James MacGlashan, Kaushik Subramanian,
Thomas J Walsh, Roberto Capobianco, Alisa Devlic, Franziska Eckert, Florian Fuchs, et al. Out-
racing champion gran turismo drivers with deep reinforcement learning. *Nature*, 602(7896):223–
228, 2022.
- Denis Yarats, Rob Fergus, Alessandro Lazaric, and Lerrel Pinto. Reinforcement learning with pro-
totypical representations. In *International Conference on Machine Learning*, pp. 11920–11931.
PMLR, 2021.

Supplementary Materials

The following content was not necessarily subject to peer review.

A Experimental Setup

A.1 Environments

In our experiments, we evaluate ProtoCRL on the *DoorKey*, *LavaCrossing*, and *SimpleCrossing* MiniGrid tasks (Chevalier-Boisvert et al., 2018). Figure 5 shows a random instantiations of each environment. The exact positions of lava, gap in the wall, door, key, and the color of the key-door combination are randomly set in each episode. The goal of the agent is to reach the green square starting from the top-left corner of the grid. The agent’s action space includes 5 actions: *forward*, *left*, *right*, *pickup key*, and *toggle door*. In each environment, the agent receives an observations consisting of:

- An egocentric 9x9 localized forward-view image (highlighted in Figure 5).
- A boolean flag indicating whether the agent is carrying an object.
- A 2D representation (*category*, *color*) of the object it is carrying. The default value is $(-1, -1)$.
- The agent’s 3D grid position and orientation (x, y, θ) .

The agent gets a reward of -0.05 for each environment step, $+1$ for reaching the goal, and -5 for falling into lava. The episode terminates when the agents reaches the goal, falls into lava, or we have reached the maximum episode length. The latter is set to 2000 while collecting experiences and 200 at evaluation time for all environments. During training, instead, the maximum episode length is set to 500 for *LavaCrossing* and *SimpleCrossing*, and to 1000 for *DoorKey*.

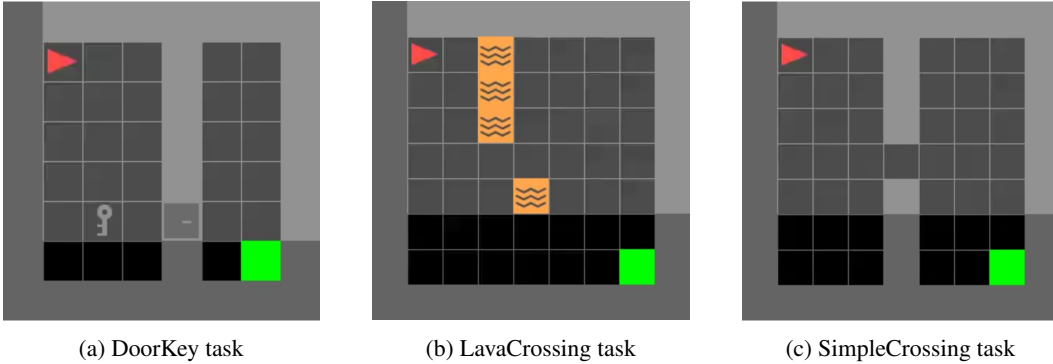


Figure 5: Random instantiations of the used MiniGrid environments.

A.2 Learning Parameters

Tables 4, 5, 7, and 6 list the parameters used in our experiments. Event conditions check whether the agent has reached the goal (*done*), has picked up a key (*pickup_key*), has opened the door (*at_door*), has fallen into lava (*at_lava*), or has traversed the gap in the wall (*at_gap*). In order to produce results for ContinualDreamer (Kessler et al., 2023), we have used the author’s original implementation. The only changes we made concerned the reward function, updated to match the one described in Section A.1, the number of environment interactions and the evaluation frequency.

Table 4: Learning parameters for experiments with no event tables (NoET).

Parameter	Value
Number of tasks (T)	3
Value function networks	3 hidden layers of 128, 128, and 64 ReLU units
Learning rate	0.0005
Batch size	64
Epsilon-greedy (ϵ)	0.3
Stale network refresh rate	0.01
Discount factor	0.99

Table 5: Learning parameters for experiments with just goal event tables (GoalET).

Parameter	Value
Number of tasks (T)	3
Event conditions (ω_i)	$done_t$ with $t = [0, T]$
Event history length (τ)	50
Event sampling probabilities (η_i)	$\mathcal{B}_0^0 : 0.3529, \mathcal{B}_1^0 : 0.1765, \mathcal{B}_2^0 : 0.2941, done_t : 0.0588$ with $t = [0, T]$
Value function networks	3 hidden layers of 128, 128, and 64 ReLU units
Learning rate	0.0005
Batch size	64
Epsilon-greedy (ϵ)	0.3
Stale network refresh rate	0.01
Discount factor	0.99

Table 6: Learning parameters for experiments with multiple event tables per task (ET).

Parameter	Value
Number of tasks (T)	3
Event conditions (ω_i)	$done_t$ with $t = [0, T]$, $pickup_key$, at_door , at_lava
Event history length (τ)	50
Event sampling probabilities (η_i)	$\mathcal{B}_0^0 : 0.34, done_0 : 0.03, at_door : 0.07, pickup_key : 0.03, \mathcal{B}_1^0 : 0.16, done_1 : 0.02, at_lava : 0.05, \mathcal{B}_2^0 : 0.23, done_2 : 0.02, at_gap : 0.05$
Value function networks	3 hidden layers of 128, 128, and 64 ReLU units
Learning rate	0.0005
Batch size	64
Epsilon-greedy (ϵ)	0.3
Stale network refresh rate	0.01
Discount factor	0.99

Table 7: Learning parameters for experiments with ProtoCRL.

Parameter	Value
Number of tasks (T)	3
Number of VGMM components (K)	12
Event conditions (ω_i)	$\mathbf{1}\left\{k = \arg \max_{j \in \{1, \dots, K\}} p(j \mid f_e(s))\right\}$
Event history length (τ)	500
Event sampling probabilities (η_i)	$\mathcal{B}_0^0 : 0.6, \mathcal{B}_1^0 : 0.3, \mathcal{B}_2^0 : 0.6, prototype_{k,t} : 0.1$ with $k \in [0, K - 1], t \in [1, T]$
Value function networks	3 hidden layers of 128, 128, and 64 LeakyReLU units
Phase 1 (\mathcal{L}_{TD}) learning rate	0.001
Phase 1 steps ratio	0.5
Phase 2 ($\mathcal{L}_{ProtoCRL}$) learning rate	0.001
Phase 2 regularization parameter (λ)	0.1
VGMM warm start epochs	10
α_{logp}	0.1
α_{KL}	0.001
β_E	1000
β_H	1.0
Batch size	64
Epsilon-greedy (ϵ)	0.3
Stale network refresh rate	0.01
Discount factor	0.99

B Continual Reinforcement Learning Metrics

In order to evaluate the CRL performance of the trained agents, we use two popular metrics: the average performance and the average forgetting.

The average performance metric estimates the average performance of the agent on all tasks at the end of the task sequence and is computed as:

$$p(t_f) = \frac{1}{T} \sum_{i=1}^T p_i(t_f)$$

where t_f is the final timestep and $p_i(t_f)$ is the performance of the agent on the i -th task at timestep t_f .

The average forgetting, instead, is measured as the average difference between the performance of each task after its own training and at the end of the task sequence:

$$F = \frac{1}{T} \sum_{i=1}^T F_i \text{ with } F_i = p_i(i \times N) - p_i(t_f)$$

where N is the number of steps per task. The forgetting for the last task in the sequence is $F_T = 0$. Finally, forgetting can also take negative values if the performance of a task i is higher at the end of the task sequence compared to after task i training.

C Episodic Return with Decreasing Buffer Capacity

Figure 6 shows the episodic return for each MiniGrid task in the tested sequence (DoorKey, LavaCrossing, SimpleCrossing) obtained by ProtoCRL and the NoET, GoalET, and ET baselines. All baselines experience severe performance drop after encountering the second task at epoch 1K, with NoET-100K/500K/1M, GoalET-100K/1M, and ET-100K/1M obtaining the lowest return of -10 at the end of the task sequence. On the other hand, ProtoCRL maintains a performance that is higher than all baselines with memory budget up to 1M transitions. Similar observations can be made for the performance on subsequent tasks, in which ProtoCRL obtains episodic returns that are slightly lower than 0 for the LavaCrossing task and slightly above 0 (highest achievable) for the SimpleCrossing task.

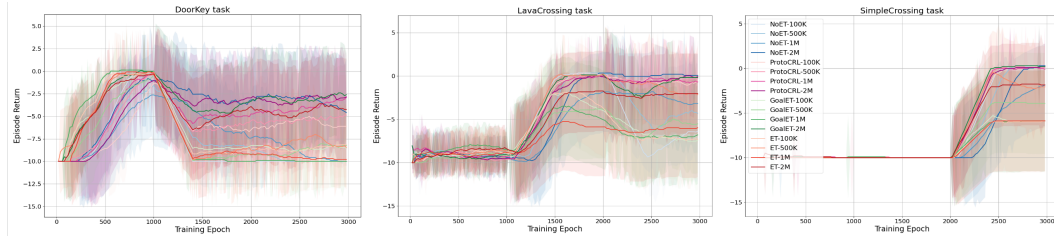


Figure 6: Episode return on 3 MiniGrid tasks using decreasing memory budget. The task ordering is *DoorKey-LavaCrossing-SimpleCrossing*. Each agent is trained on each task for 1K epochs and evaluated every 20 epochs across the entire training sequence.