

The Uniform ER Paradox: Empirical Boundaries of Memory Prioritization Under Aleatoric Uncertainty

H. M. Jisan

hjisan456@gmail.com

Islamic University

Md Mizanur Rahman

Western Sydney University

Md Shamim Hossain

University of Science and Technology of China

Research Article

Keywords:

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9760403/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

The Uniform ER Paradox: Empirical Boundaries of Memory Prioritization Under Aleatoric Uncertainty

H. M. Jisan¹, Md Mizanur Rahman^{2,3}, Md Shamim Hossain^{1,4*}

^{1*}Department of CSE, Islamic University, Kushtia, 7003, Bangladesh.

²Department of Mathematics, Islamic University, Kushtia, 7003, Bangladesh.

³School of Engineering, Design and Built Environment, Western Sydney University, Penrith, NSW, 2751, Australia.

⁴University of Science and Technology of China, Hefei, Anhui, China.

*Corresponding author(s). E-mail(s): shamim2@mail.ustc.edu.cn;

Contributing authors: hjisan456@gmail.com;

mdmizanur.rahman@westernsydney.edu.au;

Abstract

Prioritized Experience Replay (PER) is a foundational mechanism for sample-efficient Deep Reinforcement Learning (DRL). However, it suffers from severe representation collapse under aleatoric uncertainty. Specifically, unlearnable stochasticity such as the “Noisy-TV” problem permanently inflates Temporal Difference (TD) error, a phenomenon formalized here as **TD-Poisoning**, which misaligns memory sampling with actual learning progress. To systematically map these failure boundaries, we introduce Feature-Weighted Prioritized Experience Replay (FW-PER) as an experimental diagnostic probe. FW-PER strictly isolates semantic novelty by replacing TD-error in the priority distribution. The architecture utilizes an Inverse Dynamics encoder, optimized via cross-entropy, to structurally excise aleatoric noise from a low-dimensional bottleneck. This is coupled with a highly efficient $\mathcal{O}(d \log N)$ k -Nearest Neighbors global density estimator to quantify long-term structural rarity.

We present a comprehensive adversarial benchmark evaluating this mechanism across high-stakes safety (**Highway-env**), deterministic planning (**MountainCar-v0**), and extreme stochasticity (**LunarLander-v3** with Noisy-TV). Results demonstrate that FW-PER successfully acts as a semantic gatekeeper in continuous control, matching optimal safety profiles while rejecting noise,

whereas standard PER retains optimality in pure, deterministic physics. Furthermore, under extreme adversarial stochasticity, we document the “Uniform ER Paradox,” where structural filtration ultimately yields to zero-compute random sampling. Ultimately, by utilizing FW-PER to formally isolate noise, this study reveals the ‘Uniform ER Paradox,’ proving that when a state-space is overwhelmingly corrupted, the innate mathematical immunity of uniformly random sampling surpasses any algorithmic filtration.

Keywords: Deep Reinforcement Learning, Experience Replay, Aleatoric Uncertainty, Inverse Dynamics, Representation Learning, Autonomous Navigation

1 Introduction

Deep Reinforcement Learning (DRL) has achieved remarkable success in solving high-dimensional control tasks, heavily relying on Experience Replay (ER) mechanisms to break temporal correlations and improve sample efficiency [2]. This baseline uniform sampling was fundamentally advanced by Prioritized Experience Replay (PER) [5], which biases the sampling distribution toward transitions with high Temporal Difference (TD) error. By utilizing the magnitude of the TD-error as a mathematical proxy for “surprise” or expected learning progress, PER forces the agent to aggressively rehearse states where its value predictions are most inaccurate. Consequently, PER has become a ubiquitous default in modern DRL architectures.

However, the core assumption of PER is that high prediction error strictly correlates with actionable learning progress which exposes a critical vulnerability in environments characterized by aleatoric uncertainty. Biological and cognitive systems naturally filter out irreducible environmental static to focus on semantic reality; in contrast, standard prediction-based memory buffers possess no such filtering mechanism. In stochastic environments containing unlearnable noise (often conceptualized as the “Noisy-TV” problem [8]), transitions involving random elements yield a permanently high TD-error. Because neural networks cannot predict pure stochasticity, the error for these states never decays. This results in a phenomenon we formalize as **TD-Poisoning** (illustrated in Figure 1), where the replay buffer becomes systematically obsessed with unpredictable static, suffocating the sampling of controllable, high-value transitions and triggering catastrophic representation collapse [17].

To definitively diagnose and isolate this vulnerability, we introduce **Feature-Weighted Prioritized Experience Replay (FW-PER)**. Rather than relying on generative reconstruction, which often forces networks into an entropy trap attempting to model noise, FW-PER employs an Inverse Dynamics feature encoder. Optimized via cross-entropy to predict agent actions, this bottleneck naturally excises aleatoric noise, deriving a pure semantic rarity score through global k -Nearest Neighbors (k -NN) density estimation. Most critically, FW-PER performs a complete mathematical replacement of the TD-error in the priority equation, ensuring the sampling distribution is strictly insulated from TD-Poisoning.




Fig. 1 Conceptual comparison of memory prioritization under aleatoric uncertainty. (Left) Standard PER suffers from TD-Poisoning, obsessively sampling unlearnable stochasticity due to permanently high prediction error. (Right) The proposed FW-PER architecture utilizes a cognitive-inspired semantic filter, replacing TD-error to focus exclusively on structural, controllable environmental novelty.

In this paper, we do not propose FW-PER as a general-purpose algorithmic solution. Rather, we engineer FW-PER strictly as a diagnostic tool. By intentionally isolating the semantic priority signal, we use FW-PER to formally prove why standard PER collapses, and more importantly, why zero-compute Uniform ER emerges as the mathematically dominant strategy under extreme noise. Our primary contributions are:

1. **The Domain of Semantic Prioritization:** We demonstrate that in high-stakes continuous control tasks (Highway-env), pure TD-error fails to distinguish critical safety hazards from routine navigation, whereas FW-PER successfully acts as a semantic gatekeeper to reduce collisions.
2. **The Domain of TD-Error:** We prove that in clean, deterministic physics tasks (MountainCar-v0), adding semantic filtration is computationally redundant, and standard PER remains the optimal strategy.
3. **The TD-Poisoning Proof and the Baseline Paradox:** We empirically map the failure of TD-based prioritization under extreme aleatoric noise (LunarLander-v3 with Noisy TV). We demonstrate that while excising TD-error via FW-PER recovers agent stability, the optimal, most mathematically robust strategy under heavily corrupted state-spaces is to revert to zero-compute Uniform ER.

Ultimately, this study refutes the assumption of universal prioritization dominance, providing the DRL community with a definitive diagnostic map dictating when to prioritize, when to filter, and when to sample randomly based on environmental stochasticity.

2 Related Work

2.1 The Evolution and Fragility of Prioritized Replay

Experience Replay (ER) fundamentally stabilized deep reinforcement learning by decoupling the training data from the temporal stream of agent experience [2, 4]. This foundational uniform sampling was subsequently optimized by Prioritized Experience Replay (PER) [5], which demonstrated that biasing sampling distributions toward transitions with high Temporal Difference (TD) error drastically accelerates learning. However, as DRL has expanded into increasingly complex domains, the fragility of TD-based prioritization has become a critical bottleneck.

Recent literature has aggressively documented the failure modes of standard PER. [7] and [6] highlighted how prediction-error metrics destabilize when environment dynamics are non-deterministic. More recently, [19] demonstrated the degradation of standard PER in continuous control environments, proposing actor-centric prioritizations. Furthermore, [17] explicitly identified the vulnerability of PER to environmental stochasticity, proposing uncertainty-weighted mechanisms. Yet, these contemporary solutions largely retain the fundamental TD-error in their priority formulations, attempting to scale or bound it rather than excise it. We argue that under extreme aleatoric noise, scaling a poisoned metric is mathematically insufficient; true robustness requires the complete excision of the unlearnable prediction error.

2.2 Aleatoric Uncertainty and the “Noisy-TV” Problem

The vulnerability of prediction-based algorithms to stochasticity is best conceptualized through the “Noisy-TV” problem, classically defined within the subfield of intrinsic motivation [8]. When agents utilize forward prediction error as a curiosity reward, unlearnable aleatoric noise (e.g., random static on a television screen) creates a permanently high prediction loss, paralyzing exploration [10].

While this phenomenon is heavily studied in exploration literature [20, 21], its catastrophic impact on memory systems has been historically under-addressed. Because PER utilizes the exact same mathematical heuristic (prediction error) as standard curiosity modules, it inherits the exact same vulnerability. We formalize this overlap as “TD-Poisoning.” When a state-space is corrupted by aleatoric noise, the priority queue becomes hijacked, systematically over-sampling mathematically useless transitions and inducing representation collapse.

2.3 Representation Learning and the Entropy Trap

To combat representation collapse, recent memory architectures have increasingly integrated auxiliary representation learning. Generative approaches, such as World Models [11] and recent Prioritized Generative Replay frameworks [18], force the network to reconstruct the observation space. However, these methods frequently succumb to an “entropy trap.” Because generative models optimize for pixel-perfect reconstruction, they actively waste representational capacity attempting to model the irreducible noise, rendering their derived density scores computationally meaningless for memory

filtration. Similarly, while contrastive learning methods [12, 13] offer robust, self-supervised representations, they are computationally bound by large negative batch requirements, making them inefficient for dynamic replay buffer filtration.

To bypass the entropy trap without the overhead of contrastive batches, FW-PER leverages an Inverse Dynamics architecture. Pioneered by [9] for intrinsic exploration, Inverse Dynamics trains an encoder to predict the specific action a_t taken between states s_t and s_{t+1} . Because uncontrollable noise transitions independently of the agent’s physical actions, the encoder is mathematically incentivized to physically excise the noise from the latent bottleneck. Our work is the first to re-purpose this exploration mechanism as a pure, zero-TD semantic memory filter, providing a mathematically robust shield against TD-Poisoning.

3 Methodology

To systematically diagnose and resolve the representation collapse induced by aleatoric uncertainty, this section mathematically formalizes the TD-Poisoning phenomenon inherent in standard Prioritized Experience Replay (PER). Subsequently, we detail the proposed Feature-Weighted Prioritized Experience Replay (FW-PER) architecture, dissecting the Inverse Dynamics semantic bottleneck, the global latent density estimator, and the priority replacement mechanism. The complete architectural flow is illustrated in Figure 2.

3.1 Preliminaries: The TD-Poisoning Vulnerability

Standard PER [5] biases the sampling of transition $i = (s_t, a_t, r_t, s_{t+1})$ based on the magnitude of the Temporal Difference (TD) error, δ_i . For a standard Deep Q-Network (DQN) [2], the priority p_i is defined as $p_i = |\delta_i| + \epsilon$, where ϵ is a small positive constant to ensure non-zero sampling probability. The TD-error is calculated as:

$$\delta_i = r_t + \gamma \max_{a'} Q(s_{t+1}, a'; \theta_{target}) - Q(s_t, a_t; \theta_{online}) \quad (1)$$

In deterministic environments, $\lim_{t \rightarrow \infty} |\delta_i| \rightarrow 0$ as the online network θ_{online} converges to the target Q-values. However, consider a stochastic environment governed by aleatoric uncertainty (the “Noisy-TV” problem [8]). We can formalize the next state as $s_{t+1} = \hat{s}_{t+1} \oplus \xi$, where $\hat{s}_{t+1}$ represents the predictable, deterministic physical transition and $\xi \sim \mathcal{N}(0, \sigma^2)$ represents pure, unlearnable stochastic noise.

Because a neural network cannot minimize the cross-entropy or mean squared error associated with true randomness, the value prediction for s_{t+1} remains permanently inaccurate. Consequently, the expected value of the TD-error becomes bounded above zero: $\mathbb{E}[|\delta_i|] > c > 0$, where c is strictly proportional to the noise variance σ^2 . As training progresses, the priority queue becomes monopolized by these stochastic transitions, a phenomenon we term *TD-Poisoning* resulting in severe representation collapse as controllable, high-value states are systematically starved of rehearsal.




Fig. 2 The complete FW-PER Architecture. The raw observation s_t is compressed by the state encoder f_ψ into a low-dimensional bottleneck z_t . To optimize this bottleneck, an inverse model g_ω predicts the discrete action a_t taken between z_t and z_{t+1} , effectively excising unlearnable aleatoric noise. Concurrently, z_t is queried against a dynamically updated global latent bank ($\mathcal{B}_{global}$) using k -NN to compute the pure semantic rarity score ρ_i , which completely replaces TD-error in the replay buffer’s sampling distribution.

3.2 Semantic Bottleneck via Inverse Dynamics

To construct a memory filter immune to TD-Poisoning, FW-PER must extract a representation that encodes controllable physics while mathematically excising ξ . Generative models [11] fail here by attempting to reconstruct the entirety of s_{t+1} . Instead, we utilize an Inverse Dynamics architecture [9], as detailed in Figure 3.

Let f_ψ be a state encoder that maps a high-dimensional observation s_t into a low-dimensional latent vector $z_t \in \mathbb{R}^d$:

$$z_t = f_\psi(s_t) \quad (2)$$

We train an inverse model g_ω , parameterized as a Multi-Layer Perceptron (MLP), which concatenates the current latent state and the next latent state to predict the action $\hat{a}_t$:

$$\hat{a}_t = g_\omega(z_t, z_{t+1}) \quad (3)$$

For a discrete action space $|A|$, the encoder f_ψ and inverse model g_ω are jointly optimized using the cross-entropy loss between the predicted action distribution and




Fig. 3 The Inverse Dynamics feature extraction loop. By predicting a_t , the encoder is forced to discard dimensions associated with uncontrollable noise ξ .

the true one-hot action a_t :

$$\mathcal{L}_{inv}(\psi, \omega) = -\frac{1}{|B|} \sum_{i \in B} \sum_{j=1}^{|A|} a_{i,j} \log(\hat{a}_{i,j}) \quad (4)$$

Proof of Noise Excision: The aleatoric noise ξ transitions independently of the agent’s action a_t . Therefore, the mutual information $I(\xi; a_t) = 0$. Because the network g_ω is strictly optimized to minimize the prediction error of a_t , the encoder f_ψ is mathematically disincentivized from allocating latent dimensions in z_t to ξ . Consequently, the bottleneck naturally drops the unlearnable noise, isolating the semantic physics.




Fig. 4 The k -NN density estimator calculating the structural rarity of z_t within the global latent bank.

3.3 Global Latent Density Estimation

With a noise-invariant latent space established, we compute a semantic rarity score. To prevent the “in-batch blindness” common in contrastive approaches [12], FW-PER maintains a dynamic global latent bank, $\mathcal{B}_{global}$. Every M environment steps, the agent samples N_{bank} raw states uniformly from the replay buffer and passes them through the live encoder f_ψ to refresh $\mathcal{B}_{global}$.

For a sampled transition i with latent state z_i , we calculate the Euclidean distance $d(z_i, z_j) = \|z_i - z_j\|_2$ to its k -nearest neighbors within $\mathcal{B}_{global}$ (visualized in Figure 4). The local density estimate $D(z_i)$ is formulated as:

$$D(z_i) = \left(\frac{1}{k} \sum_{j=1}^k \|z_i - z_j^{(knn)}\|_2 + \eta \right)^{-1} \quad (5)$$

where η is a small constant (10^{-5}) preventing division by zero. The final semantic rarity score, ρ_i , is derived by taking the log-inverse of the density:

$$\rho_i = \ln \left(1 + \frac{1}{D(z_i) + \eta} \right) \quad (6)$$

Transitions representing rare, controllable physical states yield a high ρ_i , whereas common, highly traversed states yield a low ρ_i .

3.4 Priority Replacement and Importance Sampling

Unlike hybrid architectures that attempt to scale the TD-error by novelty ($p_i = |\delta_i| \cdot \rho_i$), FW-PER hypothesizes that the unbounded nature of aleatoric TD-error will inevitably overpower any bounded semantic penalty. Therefore, we execute a complete priority replacement. The sampling priority p_i is strictly defined as:

$$p_i = \rho_i + \epsilon \quad (7)$$

The probability of sampling a transition i from a buffer of size N is defined by the stochastic prioritization protocol [5]:

$$P(i) = \frac{p_i^\alpha}{\sum_j p_j^\alpha} \quad (8)$$

Because prioritized sampling introduces a severe bias that alters the expected value of the Q-learning updates, we must anneal this bias using Importance Sampling (IS) weights w_i . To ensure the buffer remains mathematically stable, the weights are normalized by the maximum weight in the batch:

$$w_i = \left(\frac{1}{N \cdot P(i)} \right)^\beta \frac{1}{\max_j w_j} \quad (9)$$

where β is a hyperparameter annealed from an initial value β_0 to 1.0 over the duration of training. Finally, the Q-Network is updated using the IS-weighted Mean Squared Error (MSE) objective:

$$\mathcal{L}_Q(\theta) = \frac{1}{|B|} \sum_{i \in B} w_i (\delta_i)^2 \quad (10)$$

By excising δ_i from the sampling probability $P(i)$ but retaining it in the objective function $\mathcal{L}_Q$, the agent successfully learns from the value targets while its memory rehearsal is guided exclusively by structural environmental novelty.

3.5 Algorithmic Formalization

The complete operational flow of the FW-PER architecture is synthesized in Algorithm 1.

Algorithm 1 Feature-Weighted Prioritized Experience Replay (FW-PER)

Require: Replay buffer $\mathcal{B}$, capacities N, N_{bank} , batch size B

Require: Networks $Q_\theta, f_\psi, g_\omega$, target network $Q_{\theta'}$

Require: Hyperparameters $\alpha, \beta, \gamma, k, \eta, M$

```

1: Initialize  $\mathcal{B} \leftarrow \emptyset$ , step counter  $t \leftarrow 0$ 
2: Initialize global bank  $\mathcal{B}_{global} \leftarrow \emptyset$ 
3: for episode = 1, ...,  $E$  do
4:   Observe initial state  $s_0$ 
5:   while episode not terminated do
6:     Select action  $a_t \sim \text{greedy}(Q_\theta(s_t))$  with probability  $1 - \epsilon$ 
7:     Execute  $a_t$ , observe  $r_t, s_{t+1}$ 
8:     Store transition  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{B}$  with max priority  $p_t = \max p_i$ 
9:     if  $t \pmod{M} == 0$  and  $|\mathcal{B}| \geq N_{bank}$  then
10:      Sample  $N_{bank}$  raw states uniformly from  $\mathcal{B}$ 
11:      Update  $\mathcal{B}_{global} \leftarrow \{f_\psi(s) \mid s \in \text{sampled states}\}$ 
12:    end if
13:    if  $|\mathcal{B}| \geq B$  then
14:      Sample batch  $b \sim P(i) = p_i^\alpha / \sum p_j^\alpha$ 
15:      Compute IS weights  $w_i$  for  $i \in b$ 
16:      Compute latent features  $z_i = f_\psi(s_i), z_{i+1} = f_\psi(s_{i+1})$ 
17:      Update Inverse Models  $(f_\psi, g_\omega)$  via  $\nabla \mathcal{L}_{inv}$ 
18:      Update Q-Network  $Q_\theta$  via  $\nabla \mathcal{L}_Q$  with  $w_i$ 
19:      Compute new density  $D(z_i)$  via  $k$ -NN against  $\mathcal{B}_{global}$ 
20:      Update priorities  $p_i \leftarrow \ln(1 + (D(z_i) + \eta)^{-1}) + \epsilon$ 
21:    end if
22:     $t \leftarrow t + 1$ 
23:  end while
24: end for

```

4 Experimental Setup

To rigorously map the operating boundaries of experience replay mechanisms under varying degrees of aleatoric uncertainty and reward sparsity, we design a three-tier adversarial benchmark. We evaluate our proposed Feature-Weighted Prioritized Experience Replay (FW-PER) against three baselines: Uniform Experience Replay (Uniform ER) [2] representing the zero-compute, unbiased sampling paradigm; Standard Prioritized Experience Replay (PER) [5] representing the foundational prediction-error paradigm; and Uncertainty Prioritized Experience Replay (UPER) [17] representing the contemporary state-of-the-art epistemic uncertainty approach via neural ensembles. To ensure statistical significance and account for initialization variance, all experiments are executed across five independent random seeds (42, 100, 314, 1234, 9999), with performance metrics smoothed using a moving average window and shaded regions representing the standard deviation.

4.1 Baseline Selection Rationale

By bridging foundational architectures (Uniform ER, PER) with modern ensemble-based uncertainty estimators (UPER), we mathematically isolate the exact threshold where pure prediction-error fails under aleatoric noise. Furthermore, comparing our single-network $\mathcal{O}(1)$ FW-PER against the computationally heavy $\mathcal{O}(N)$ UPER ensemble allows us to strictly evaluate whether tracking environmental variance (aleatoric) offers a more computationally efficient survival mechanism than tracking network variance (epistemic).

4.2 Hardware and Implementation Details

All empirical evaluations, environment simulations, and neural network training phases were executed using Python and the PyTorch framework. The computational pipeline was deployed on an NVIDIA T4 Tensor Core GPU (16 GB VRAM), provisioned via the Google Colaboratory cloud environment. To ensure strict statistical reproducibility across the independent training evaluations, all environment stochasticity, action sampling, and network weight initializations were globally seeded ($\{42, 100, 314, 1234, 9999\}$).

4.3 Network Architectures and Parameterization

To isolate the specific impact of the memory buffer mechanisms, the core agent across all experiments is standardized as a Deep Q-Network (DQN) [2]. The online value network Q_θ and the target network $Q_{\theta'}$ are parameterized as identical Multi-Layer Perceptrons (MLPs) consisting of an input layer corresponding to the environment’s observation space, two hidden layers of 256 units each utilizing Rectified Linear Unit (ReLU) activations, and a linear output layer corresponding to the discrete action space $|A|$.

For the FW-PER architecture, the Inverse Dynamics module is similarly parameterized to ensure computational fairness. The state encoder f_ψ processes the raw observation through a single 64-unit ReLU hidden layer, projecting it into a strictly

low-dimensional latent bottleneck of $d = 4$. The inverse model g_ω concatenates the current and subsequent latent states ($[z_t, z_{t+1}] \in \mathbb{R}^8$) and processes them through a 64-unit ReLU hidden layer to output the predicted action logits. Both the DQN and the Inverse Dynamics networks are optimized using the Adam optimizer [16]. The complete hyperparameter suite, optimized via preliminary grid search, is documented in Table 1.

Table 1 Comprehensive Hyperparameter Configuration for Reproducibility

Module	Hyperparameter	Value
DQN Optimization	Learning Rate (α_{DQN})	5×10^{-4}
	Optimizer	Adam
	Discount Factor (γ)	0.99
	Target Update Frequency (M_{target})	1000 steps
	Exploration Fraction (ϵ -decay)	0.995
	Minimum Exploration (ϵ_{min})	0.05
Inverse Dynamics	Encoder Latent Dimension (d)	4
	Inverse Model Learning Rate (α_{inv})	1×10^{-3}
	Density Neighbors (k)	10
	Density Constant (η)	1×10^{-5}
Global Latent Bank	Bank Capacity (N_{bank})	2000
	Refresh Frequency ($M_{refresh}$)	1000 steps
	Bank Sampling Strategy	Uniform
Replay Mechanisms	Buffer Capacity (N)	15000
	Batch Size (B)	128
	Priority Exponent (α)	0.6
	IS Weight Annealing (β)	$0.4 \rightarrow 1.0$
	Priority Constant (ϵ)	1×10^{-3}

4.4 Domain I: Semantic Prioritization in Continuous Control

The first environment, *Highway-Fast-v0* [15], evaluates the algorithms in a high-stakes, multi-agent continuous control setting. The ego-vehicle navigates a dense traffic simulation, possessing a discrete meta-action space (e.g., change lane left/right, accelerate, decelerate, idle). The observation space is a 5×5 kinematic matrix encoding the position and velocity of the ego-vehicle and the four nearest surrounding vehicles. The reward function heavily penalizes collisions ($R_{collision} = -1.0$) while rewarding high-speed traversal. In this domain, standard TD-error frequently fails to differentiate between high-error states caused by routine lane-change miscalculations and critical safety failures. This environment tests FW-PER’s capacity to act as a semantic gatekeeper, prioritizing the structurally rare pre-collision states to enforce safer driving policies.

4.5 Domain II: The Deterministic Physics Baseline

To establish the lower operational boundary of the proposed architecture, we deploy the agents in *MountainCar-v0* [14]. This classic control environment features a perfectly deterministic, 2-dimensional continuous state space (position $x \in [-1.2, 0.6]$ and velocity $v \in [-0.07, 0.07]$) and a discrete action space of 3 (push left, push right, no push). The reward is strictly sparse: $R_t = -1$ for every time step until the agent reaches the goal flag at $x \geq 0.5$. Because this environment contains zero aleatoric noise, the standard TD-error provides a mathematically pure signal of expected learning progress. This domain explicitly tests whether the Inverse Dynamics and k -NN density calculations in FW-PER introduce unnecessary computational overhead and signal dilution in environments where semantic filtration is objectively redundant.

4.6 Domain III: Adversarial Aleatoric Noise (The “Noisy-TV”)

The core of our benchmark tests the vulnerability of priority mechanisms to unlearnable stochasticity. We utilize *LunarLander-v3* [14], a complex 8-dimensional physics simulation requiring the agent to safely land a spacecraft. The base observation $s_{base} \in \mathbb{R}^8$ includes coordinates, velocities, angle, angular velocity, and Boolean leg contacts.

To induce the TD-Poisoning phenomenon, we construct an adversarial “Noisy-TV” wrapper. At every discrete time step t , the environment generates a high-variance, 10-dimensional aleatoric noise vector ξ_t , where each dimension is sampled from a Gaussian distribution:

$$\xi_t \sim \mathcal{N}(0, 5) \quad \text{with} \quad \xi_t \in \mathbb{R}^{10} \quad (11)$$

This noise vector is concatenated to the base observation, creating a corrupted 18-dimensional state space:

$$s'_t = s_{base,t} \oplus \xi_t \in \mathbb{R}^{18} \quad (12)$$

Because ξ_t is highly variable and completely independent of the agent’s actions, the DQN’s value predictions for s'_{t+1} yield a permanently high TD-error. This environment acts as an adversarial stress test, designed to purposefully trigger catastrophic representation collapse in standard PER by hijacking its sampling distribution, thereby proving the necessity of FW-PER’s semantic priority replacement mechanism.

5 Results and Discussion

We present the empirical results of our adversarial benchmark across the three defined domains. The performance of Feature-Weighted Prioritized Experience Replay (FW-PER) is evaluated against Standard PER and Uniform ER. Each configuration was trained across 5 independent random seeds, with the solid lines representing the mean performance and the shaded regions representing the standard deviation ($\pm 1\sigma$).

Highway-Fast-v0: Multi-Seed Statistical Performance




Fig. 5 Multi-seed statistical performance in the Highway-Fast-v0 domain. The graphs denote (Left) Sample Efficiency via smoothed reward, (Center) Cumulative crashes indicating the safety profile, and (Right) Average driving speed indicating navigation competence.

5.1 Domain I: Semantic Prioritization in Continuous Control

The *Highway-Fast-v0* environment serves as our testbed for high-stakes semantic filtering. As depicted in Figure 5, standard prediction-error metrics face a unique challenge in this domain: both highly successful, high-speed navigational maneuvers and catastrophic collisions generate massive temporal difference (TD) errors.

The empirical data highlights a critical divergence in agent behavior. Uniform ER demonstrates aggressive sample efficiency early in training but achieves this by adopting a highly reckless policy, evidenced by the highest cumulative collision rate (Figure 5, Center) and a severe degradation in average speed at the tail end of training (Figure 5, Right) as the agent learns to overly idle to avoid crashes.

Conversely, FW-PER and Standard PER successfully suppress the collision rate. However, standard PER achieves safety by sacrificing driving competence, dropping its average speed significantly in later episodes. FW-PER emerges as the optimal semantic gatekeeper. By utilizing the Inverse Dynamics semantic rarity score (ρ_i) rather than raw TD-error, FW-PER accurately prioritizes structurally rare, near-miss states without obsessing over the mathematical shock of terminal collisions. Consequently, FW-PER maintains the highest sustained average speed (~ 25 m/s) while keeping the cumulative crash rate statistically tied with the safest baseline, proving that semantic filtering successfully balances high-competence navigation with critical safety constraints.

5.2 Domain II: The Redundancy of Semantics in Deterministic Physics

To establish the mathematical ceiling of standard prediction-based memory, we evaluate the architectures in the purely deterministic *MountainCar-v0* environment (Figure 6). Because this environment is devoid of aleatoric noise, the TD-error acts as an uncorrupted, perfectly pristine signal of expected learning progress.




Fig. 6 Performance in the MountainCar-v0 sparse-reward benchmark. (Left) Smoothed reward tracking the agent’s ability to escape the valley. (Right) The success rate percentage of reaching the goal state across 5 seeds.

The results definitively confirm our initial hypothesis regarding the operational boundaries of semantic filtering. Standard PER dominates this domain, achieving a 100% solve rate significantly faster than all alternatives and converging to an asymptotic reward ceiling of roughly -120 . Uniform ER completely collapses under the sparse-reward constraint, failing to consistently solve the environment, underscoring why PER was originally invented.

Crucially, while FW-PER successfully solves the environment, it fundamentally lags behind standard PER. This empirical lag is an intentional diagnostic finding. It proves our hypothesis that introducing semantic bottlenecks in clean physics actively destroys pristine TD signals. This empirical gap proves that in clean, low-dimensional deterministic physics, introducing an Inverse Dynamics bottleneck is not only computationally redundant but actively dilutes the purity of the TD-error signal. This transparent ablation solidifies our claim that FW-PER is a specialized surgical tool for uncertainty, not a universal replacement for clean physics engines.

5.3 Domain III: TD-Poisoning and the Baseline Paradox

The climax of our adversarial benchmark evaluates the memory systems under extreme unlearnable stochasticity. Figure 7 illustrates the catastrophic failure of prediction-based sampling and validates the proposed semantic priority replacement mechanism.

As formalized in Section 3, standard PER suffers from severe TD-Poisoning. Because the 10-dimensional Gaussian noise injected into the state space is unlearnable, the online Q-network’s predictions are permanently inaccurate. The priority queue becomes hyper-fixated on this static. Figure 7 (Left) captures this mathematical collapse perfectly: Standard PER completely flatlines around a reward of -150 and exhibits a heavily degraded survival duration (Figure 7, Right).

By executing a complete mathematical replacement of the TD-error and relying strictly on the Inverse Dynamics rarity score, FW-PER successfully insulates the agent from representation collapse. Because the Inverse Model is optimized via cross-entropy

LunarLander-v3: Robustness to Aleatoric Noise




Fig. 7 Adversarial stress test using LunarLander-v3 with a Noisy-TV wrapper ($\mathcal{N}(0, 5.0)$). (Left) Sample efficiency and representation collapse. (Center) Task solve percentage. (Right) Steps survived per episode, measuring policy stability under extreme noise.

to predict controllable actions, it successfully drops the aleatoric noise from the latent bottleneck. Consequently, FW-PER effectively ignores the Noisy-TV, detaches from the Standard PER failure curve, and continuously climbs toward the zero-reward threshold.

The Uniform ER Paradox: The most profound finding of this benchmark is the performance of Uniform ER. Under extreme aleatoric noise, Uniform ER achieves the highest sample efficiency and highest survival rate. This presents a critical paradox in modern DRL: because Uniform ER samples transition data completely randomly, it possesses an innate, zero-compute immunity to TD-Poisoning. While FW-PER successfully repairs the vulnerability of prioritized sampling through sophisticated filtering, this empirical data proves that when a state-space is overwhelmingly corrupted by irreducible noise, no mathematical filter surpasses the raw robustness of uniformly random sampling.

5.4 Ablation Study: Hyperparameter Sensitivity

To ensure the performance gains observed in FW-PER are derived from the core architectural methodology rather than hyperparameter over-fitting, we conducted an ablation study focusing on the two most critical parameters of the semantic filter: the latent bottleneck dimension (d) and the nearest-neighbor density parameter (k).

Latent Bottleneck Dimension (d): The dimensionality of the Inverse Dynamics encoder dictates the strictness of the semantic filter. We evaluated FW-PER on the Noisy-TV benchmark using $d \in \{2, 4, 16, 64\}$. Empirical results demonstrated that a large latent space ($d \geq 16$) induces a regression toward standard PER failure. With excess capacity, the network is no longer mathematically forced to drop the aleatoric noise ξ_t , allowing stochasticity to bleed back into the density calculation and trigger TD-Poisoning. Conversely, an overly constrained bottleneck ($d = 2$) resulted in severe physical representation collapse, where the agent could not distinguish between distinct, controllable states. The $d = 4$ bottleneck proved optimal, providing just enough capacity to encode physical coordinates while strictly rejecting high-variance noise.

Density Estimator Resolution (k): We further analyzed the sensitivity of the k -NN global density estimator by varying $k \in \{1, 5, 10, 50\}$. For $k = 1$, the priority distribution became highly unstable, as the rarity score ρ_i was overly sensitive to outlier states in the global bank $\mathcal{B}_{global}$. As k approached 50, the density scores homogenized, resulting in a sampling distribution that effectively mimicked Uniform ER. Setting $k = 10$ provided the optimal balance, yielding a smooth, stable semantic rarity curve that accurately identified structural novelty without being derailed by localized latent anomalies.

5.5 Computational Complexity and Overhead Analysis

A primary critique of advanced replay mechanisms is the introduction of prohibitive computational overhead. Standard Uniform ER operates with an update and sampling complexity of $\mathcal{O}(1)$. Standard PER [5] utilizes a SumTree data structure, increasing the time complexity for both sampling and updating to $\mathcal{O}(\log N)$, where N is the buffer capacity.

The proposed FW-PER architecture introduces two additional computational pathways: the Inverse Dynamics forward pass and the k -NN density estimation. Let $|W|$ denote the number of parameters in the state encoder f_ψ . The computation of the latent vector z_i requires $\mathcal{O}(|W|)$ operations. However, because z_i is exceptionally low-dimensional ($d = 4$), the subsequent density estimation is highly optimized.

By structuring the global latent bank $\mathcal{B}_{global}$ as a KD-Tree, the k -NN search complexity is reduced to $\mathcal{O}(d \log N_{bank})$. Because the bank size ($N_{bank} = 2000$) is significantly smaller than the total buffer capacity ($N = 15000$), and the bank is only refreshed every $M = 1000$ environment steps, the amortized computational overhead of the semantic filter is minimal. During our empirical evaluations, FW-PER increased the wall-clock training time by only 14.2% compared to standard PER. This proves that the semantic filter introduces minimal overhead, rendering FW-PER a highly viable, lightweight solution for stochastic domains.

6 Limitations and Real-World Applicability

While FW-PER establishes a rigorous defense against TD-Poisoning, we acknowledge several architectural limitations that define vectors for future research.

First, the current architecture relies on a static global latent bank ($\mathcal{B}_{global}$) that is refreshed at a hard-coded frequency (M). In highly dynamic, non-stationary environments, a periodic refresh may result in "latent staleness," where the semantic rarity scores ρ_i are computed against an outdated representation of the agent's behavioral distribution. Future iterations should explore dynamic, surprise-triggered refresh schedules.

Second, the Inverse Dynamics bottleneck inherently assumes that the controllable elements of the environment are fully sufficient to maximize the reward function. In environments where the aleatoric noise itself is tied to the reward signal (e.g., predicting stochastic financial market fluctuations), excising the noise from the latent

representation would actively blind the agent to critical task dynamics. Therefore, FW-PER is strictly applicable to environments where stochasticity acts as an adversarial distractor rather than a core mechanic of the objective function.

Finally, while FW-PER maps the boundary where Uniform ER outperforms mathematically filtered buffers under extreme noise, it requires a manual architectural switch. A truly autonomous, human-like cognitive architecture would possess a meta-controller capable of tracking the variance of its own TD-error in real-time, dynamically interpolating between Uniform ER, PER, and FW-PER as the environmental stochasticity shifts during deployment.

7 Conclusion and Future Work

The ubiquitous deployment of Prioritized Experience Replay (PER) in Deep Reinforcement Learning (DRL) relies on the foundational assumption that Temporal Difference (TD) error serves as a universally optimal proxy for learning progress. In this paper, we systematically dismantled this assumption, demonstrating that under conditions of aleatoric uncertainty, prediction-based prioritization algorithms suffer from catastrophic representation collapse. By formalizing this vulnerability as “TD-Poisoning,” we proved that unlearnable stochasticity such as the Noisy-TV problem permanently inflates the prediction error, causing the replay buffer to obsessively over-sample mathematically useless noise at the expense of controllable, high-value transitions.

To mathematically dissect this failure mode without falling into the generative entropy trap, we utilized Feature-Weighted Prioritized Experience Replay (FW-PER) as a diagnostic baseline. Drawing inspiration from cognitive systems capable of filtering environmental static, FW-PER strictly isolates semantic reality by replacing the poisoned TD-error in the priority distribution. By utilizing an Inverse Dynamics encoder optimized via cross-entropy, the architecture physically excises aleatoric noise from a low-dimensional bottleneck. This noise-invariant representation is then coupled with a computationally efficient $\mathcal{O}(d \log N_{bank})$ global k -Nearest Neighbors density estimator to quantify structural, long-term environmental novelty.

Our comprehensive adversarial benchmark established a definitive operational boundary map for memory prioritization mechanisms:

1. **Continuous Control and Safety:** In high-stakes, multi-agent navigation (Highway-Fast-v0), FW-PER successfully functioned as a semantic gatekeeper, ignoring the chaotic TD-variance of standard driving to prioritize structurally rare pre-collision states, thereby matching the safest baselines while maintaining superior driving competence.
2. **Deterministic Redundancy:** In pure, sparse-reward physics (MountainCar-v0), standard PER remained mathematically optimal. Our ablation confirmed that deploying semantic filtration in environments devoid of noise introduces unnecessary computational overhead and signal dilution, proving that FW-PER is a specialized surgical tool rather than a universal heuristic.
3. **The Baseline Paradox:** Under extreme adversarial stochasticity (LunarLander-v3 with Noisy-TV), standard PER suffered complete representational collapse.

While FW-PER successfully excised the injected noise and recovered learning stability, the zero-compute Uniform Experience Replay ultimately achieved the highest sample efficiency. This profound “Uniform ER Paradox” establishes that when a state-space is overwhelmingly corrupted, the innate mathematical immunity of uniformly random sampling surpasses any algorithmic filtration.

Future Work: While FW-PER provides a robust mathematical shield against TD-Poisoning, the ultimate trajectory of this research points toward biologically plausible, autonomous memory systems. Biological cognition does not rely on a single, static memory filter; it dynamically adapts its attention based on the perceived entropy of the environment. Future work will focus on the development of an active meta-controller capable of tracking the variance of its own TD-error in real-time. By estimating the environmental stochasticity ratio dynamically, such an agent could automatically toggle between pure TD-error (for deterministic planning), semantic filtration (for structural exploration), and uniform sampling (for pure noise survival) without manual architectural intervention. Achieving this dynamic routing will represent a critical step toward realizing truly autonomous, human-like cognitive architectures in Deep Reinforcement Learning.

Declaration of competing interest

The authors declare that there are no financial conflicts of interest or personal relationships that may have impacted the work presented in this paper.

Credit authorship contribution statement

H. M. Jisan: Writing – Original Draft, Methodology, Formal Analysis, Conceptualization, Visualization, Validation. **Md Mizanur Rahman:** Writing – Review & Editing. **Md Shamim Hossain:** Writing – Review & Editing, Supervision.

Data, Code, and Availability

To ensure the absolute reproducibility of this micro-regime ablation study, all code, hyperparameters, and resulting datasets have been made publicly available.

- **Source Code:** The complete PyTorch implementation of the FW-PER architecture, the comparison models and the other architectures are hosted on GitHub at: <https://github.com/hmjisan/FW-PER>.
- **Execution Environment:** All experiments were executed sequentially utilizing a cloud-based virtual machine provisioned with a single NVIDIA T4 Tensor Core GPU (16GB VRAM). The repository includes the exact Jupyter notebooks required to replicate the hardware constraints.

Acknowledgments

The authors would like to thank the Computer Vision and NLP lab and the Department of Computer Science and Engineering at Islamic University, Bangladesh, for

providing the research environment and facilities necessary to conduct this work. We also gratefully acknowledge the creators of the open-source environments used in this study.

References

- [1] Sutton RS, Barto AG (2018) Reinforcement learning: An introduction. MIT press.
- [2] Mnih V, Kavukcuoglu K, Silver D, et al. (2015) Human-level control through deep reinforcement learning. *Nature* 518(7540):529–533.
- [3] Van Hasselt H, Guez A, Silver D (2016) Deep reinforcement learning with double q-learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol 30, no 1.
- [4] Lin LJ (1992) Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning* 8(3):293–321.
- [5] Schaul T, Quan J, Antonoglou I, Silver D (2015) Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- [6] Horgan D, Quan J, Budden D, et al. (2018) Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*.
- [7] Fedus W, Ramachandran P, Agarwal R, et al. (2020) Revisiting fundamentals of experience replay. In: *International Conference on Machine Learning*, PMLR, pp 3061–3071.
- [8] Burda Y, Edwards H, Pathak D, et al. (2018) Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*.
- [9] Pathak D, Agrawal P, Efros AA, Darrell T (2017) Curiosity-driven exploration by self-supervised prediction. In: *International Conference on Machine Learning*, PMLR, pp 2778–2787.
- [10] Bellemare M, Srinivasan S, Ostrovski G, et al. (2016) Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems* 29.
- [11] Ha D, Schmidhuber J (2018) World models. *arXiv preprint arXiv:1803.10122*.
- [12] Laskin M, Srinivas A, Abbeel P (2020) Curl: Contrastive unsupervised representations for reinforcement learning. In: *International Conference on Machine Learning*, PMLR, pp 5639–5650.
- [13] Schwarzer M, Anand A, Goel R, et al. (2020) Data-efficient reinforcement learning with self-predictive representations. *arXiv preprint arXiv:2007.05929*.

- [14] Brockman G, Cheung V, Pettersson L, et al. (2016) Openai gym. arXiv preprint arXiv:1606.01540.
- [15] Leurent E (2018) An environment for autonomous driving decision-making. GitHub repository.
- [16] Kingma DP, Ba J (2014) Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
- [17] Carrasco-Davis R, Lee S, Clopath C, Dabney W (2025) Uncertainty prioritized experience replay. Reinforcement Learning Journal.
- [18] Lu C, et al. (2025) Prioritized generative replay. International Conference on Learning Representations (ICLR).
- [19] Saglam B, et al. (2023) Actor prioritized experience replay. Journal of Artificial Intelligence Research 78:639–672.
- [20] Nguyen MH, et al. (2024) Variable-agnostic causal exploration for reinforcement learning. arXiv preprint arXiv:2407.12437.
- [21] Lee S, et al. (2024) A temporally correlated latent exploration for reinforcement learning. International Conference on Learning Representations (ICLR).