Prioritized Experience Replay in Deep Q-Networks: Accelerating Learning in Atari Environments

Ayush Singh Bhadoria
Department of Computer Science Engineering, Gautam Buddha University, Greater Noida, Uttar Pradesh, India

Dr. Nitesh Singh Bhati
Department of Computer Science Engineering, Gautam Buddha University, Greater Noida, Uttar Pradesh, India

Shubham Yash Tomar
Technical Lead, Jarvis Technology Postgraduate, (IITJ), Delhi, India

Abstract:

Deep Q-Networks (DQN) have demonstrated remarkable success in learning control policies from raw pixel observations in Atari games. However, uniform sampling in experience replay—a key component of DQN—limits data efficiency by overlooking the relative importance of individual transitions. This work integrates Prioritized Experience Replay (PER) into the DQN framework, prioritizing transitions with high temporaldifference (TD) error in order to accelerate learning. We provide extensive implementation details, including a computationally efficient sum-tree structure and bias correction through importance sampling. Our experiments on several Atari 2600 games show that PER-DQN can lead to faster convergence (often by a factor of 1.8×) and significantly higher final scores (on the order of 22% average improvement) when compared to the original DQN with uniform replay. We present ablation studies revealing hyperparameter sensitivities for the PER algorithm and demonstrate consistent improvements across different settings of αand β. Additional code snippets, training curves, and architectural diagrams are offered to encourage reproducibility.

Keywords:

Deep Reinforcement Learning, Q-Learning, Experience Replay, Atari 2600, Prioritized Sampling, Neural Networks.