






Learning from **connected** n -step transitions with expectile loss can be as efficient as a large replay buffer.



Endpoint Replay: Compressing the Recency Buffer in Deep Reinforcement Learning

Parham Mohammad Panahi, Armin Ashrafi, Haoyu Du, Andrew Patterson, Martha White, Adam White



UNIVERSITY
OF ALBERTA