




Squeezing More from the Stream


Learning Representation Online for Streaming Reinforcement Learning

Nilaksh, Antoine Clavaud, Mathieu
Reymond, François Rivest, Sarath
Chandar


[nilaksh.nilaksh@mila.quebec](mailto:nilaksh.Nilaksh@mila.quebec)

What is the streaming RL setting?

	Replay Buffer	Batch Updates	Parallelized Environments
Batch RL	yes	yes	yes
Streaming RL	no	no	no




(a) Batch RL




(b) Streaming RL


Adding Representation Loss to the Streaming RL Pipeline



Aggregate Scores on Three Benchmarks



Performance profiles




Latent Space Analysis

Effective Encoder ranks: Higher is better

Algorithm	Without SPR	With SPR (ours)
QRC(λ)	42.7 $\pm$ 33.7	1256.9 $\pm$ 296.5
Stream Q(λ)	466.4 $\pm$ 120.4	912.5 $\pm$ 209.5
Streaming DQN	518.7 $\pm$ 259.4	526.7 $\pm$ 243.4

With
SPR



No
SPR

