



Investigating Prioritized Replay and how it interacts with Generalization via Neural Nets.


Insight 1: Prioritization improves sample efficiency over uniform experience replay in tabular Markov chain but it **overestimates** when combined with neural nets.




Insight 2: Prioritized updates over-generalize when using TD-learning with neural networks.




Insight 3: Updating priorities can improve learning speed in tabular control.



BUT this finding **does not** extend to the neural network setting..



Problem Setting:
Prediction and Control in Markov Chain



See paper for more insights!