

Optimal Sparsification and Low-Rank Decomposition of Deep Reinforcement Learning Policies for Surgical Robot Task Automation

Goddla, Vikram (School: Detroit Country Day School)

Deep reinforcement learning (DRL) has shown significant promise in a wide range of applications including computer games and robotics. Yet, DRL policies require very long training times resulting in dense policies with excessive number of network connections. These dense DRL policies are prone to overfitting and consume massive computing resources limiting their performance in real-world applications such as robotic surgery. Pruning, and singular-value-decomposition have been proposed to achieve model sparsification and compression to limit overfitting and reduce memory consumption. However, they resulted in sub-optimal performance with significant decay in rewards. L1 and L2 regularization have been proposed for sparsification in autoencoders, but they do not induce optimal sparsity due to their bias toward larger weights and their implementation in DRL is not apparent. I developed a novel L0-regularization framework using an optimal sparsity map to sparsify DRL policies and promote their decomposition to a lower rank without decay in rewards. I evaluated my L0-regularization framework across five environments (Cartpole, Acrobat, LunarLander, SuperMarioBros and Surgical Robot Learning) using several on-policy and off-policy algorithms. I demonstrated that the L0-regularized DRL policy in the SuperMarioBros environment achieved 93% sparsity and gained 70% compression while outperforming the dense policy by 200 points. Additionally, the L0-regularized DRL policy in the Surgical Robot Learning environment, achieved a 36% sparsification and gained 46% compression, while significantly outperforming dense policies, providing solid evidence that L0-regularized sparse DRL policies are very effective in limiting overfitting and reducing computational resource requirements.