REPNET: A Reproductive Training Architecture for Reinforcement Learning Neural Networks

Lindauer, Tiernan (School: Vista Ridge High School)

Since 2001, actor-critic algorithms have combined “the strong points of actor-only and critic- only methods” (Konda & Tsitsiklis, 1999) to reduce the time needed to train a reinforcement learning neural network. The REProductive NETwork architecture, REPNET, uses a custom data structure built on top of this idea to train networks in significantly less time. During training, the models in REPNET (called branches) can reproduce, rewarding models that achieve a high running reward change. Branches can also be pruned due to a lack of reproduction, penalizing stagnant branches. These branches are contained in the custom recursive data structure of REPNET. Through the use of automatic add-ons such as Adaptive Pruning Time Adjustment (APTA) and Auto-Generating Nonlinear Threshold Functions, these hyperparameters can be tuned without human input. After training the neural network on OpenAI’s “CartPole” environment, REPNET reduced the mean of the training episodes by 29.3% and the standard deviation of episodes by 72.62%. Overall, REPNET made reinforcement learning training significantly faster and less variable in comparison to its actor-critic counterpart. Decreasing the training times of reinforcement learning neural networks can accelerate progress in AI research, leading to faster development of technologies that can improve various aspects of the world, from healthcare to environmental sustainability.