The Lottery Ticket Hypothesis and value-based Deep Reinforcement Learning

Auteur : Debes, Baptiste
Promoteur(s) : Louppe, Gilles
Faculté : Faculté des Sciences appliquées
Diplôme : Master : ingénieur civil en science des données, à finalité spécialisée
Année académique : 2021-2022
URI/URL : http://hdl.handle.net/2268.2/13872

University of Liège
Faculty of Applied Sciences
Department of Electrical Engineering and Computer Science

The Lottery Ticket Hypothesis and value-based Deep Reinforcement Learning

A dissertation submitted in partial fulfillment of the requirements for the degree of
Master of Science in Data Science and Engineering

by

BAPTISTE DEBES

Supervisor:
Prof. Gilles Louppe

Academic Year 2021-2022

Abstract

The Lottery Ticket Hypothesis (LTH) suggests that randomly initialized overparametrized neural networks contain subnetworks which - when trained in isolation - are able to perform better than similar subnetworks whose architecture and weights are drawn randomly. Subnetworks matching the Lottery Ticket Hypothesis are referred to as winning tickets because they are the winners of the initialization lottery. An algorithm called Iterative Magnitude Pruning (IMP) was introduced to discover winning tickets. Finding well-performing sparse neural networks is especially interesting because of the potential large reduction in memory footprint and global computational burden. These combined may lead to an important reduction of the energy required to perform a same task. Deep Reinforcement Learning (DRL) has introduced algorithm capable of solving complex tasks (dynamic system control, Atari games, board games, ...). In this work we study the combination of deep reinforcement learning and the lottery ticket hypothesis. We focus on two algorithms namely Double Deep Q-Networks (DDQN) and Soft-Actor-Critic (SAC) which both belong to the fruitful class of value-based methods. We provide the third independent confirmation - in the context of deep reinforcement learning - of the existence of subnetworks matching the Lottery Ticket Hypothesis using Iterative Magnitude Pruning. Our experiments were carried on standard classic control as well as pixel-based environments. We provide experiments and guidelines regarding some important hyperparameters. We suggest a potential ability of winning tickets to robustly preserve low rank embeddings of the environment's state space. Some of ours results suggest that tickets found using IMP seem closer than expected to subnetworks that could be found using so-called structured pruning methods. Our experiments also showcase the ability of winning tickets to render inactive useless input variables while keeping good performance on the task. This result along with others indicate a potential ability of winning tickets to be used as feature importance extractors. Finally, a variant of Iterative Magnitude Pruning is introduced which we call pooled pruning. We suggest this variant could be beneficial for multi-networks algorithms such as Soft-Actor-Critic.