Winning the Lottery by Preserving Network Training Dynamics With Concrete Ticket Search

Arora, Tanay (School: Jesuit High School)

The Lottery Ticket Hypothesis posits that highly sparse, trainable subnetworks ('winning tickets') exist within all randomly initialized neural networks. However, state-of-the-art ticket-drawing methods, like Lottery Ticket Rewinding (LTR), are computationally prohibitive, while more efficient saliency-based Pruning-at-Initialization (PaI) techniques suffer from a significant performances drops and fail basic sanity checks. In this work, we argue that PaI's reliance on first-order saliency metrics, which ignore inter-weight dependencies, contributes substantially to this performance gap. To address this, we introduce Concrete Ticket Search (CTS), which frames subnetwork discovery as holistic combinatorial optimization. By leveraging a Concrete relaxation of the discrete search space and a novel gradient balancing scheme (GRADBALANCE) to control sparsity, CTS efficiently identifies high-performing subnetworks near initialization without requiring sensitive hyperparameter tuning. Motivated by recent works on lottery ticket training dynamics, we further propose a knowledge distillation-inspired family of pruning objectives, finding that minimizing the reverse Kullback-Leibler divergence between sparse and dense network outputs (CTS-KL) is particularly effective. Experiments on varying baseline image tasks show that CTS produces subnetworks that robustly pass sanity checks and perform close to or better than LTR, while requiring only a small fraction of the computation. For example, on ResNet-20/CIFAR10 at 99.3% sparsity, it attains 74.0% accuracy in 7.9 minutes, while LTR achieves 68.3% accuracy in 95.2 minutes. CTS's subnetworks outperform saliency methods across all sparsities, but its accuracy advantage over LTR is most pronounced in the highly sparse regime.

Awards Won: