Finding Dense Supermasks in Randomly Initialized Neural Networks

Csanád Sándor[a,b]

[a]Faculty of Mathematics and Informatics,
Babeş-Bolyai University, Cluj-Napoca, Romania
csanad.sandor@cs.ubbcluj.ro

[b]Robert Bosch SRL, Cluj-Napoca, Romania

Abstract

Recent works on network pruning [1, 2] showed that randomly weighted neural networks contain supermasks (subnetworks) the performance of which is comparable with similarly sized, trained networks. These results underpin the Lottery Ticket Hypothesis [3]: the effectiveness of deep neural networks rely on lucky initialization.

While these works define sparse supermasks, we demonstrate that dense supermasks can also be found by applying structured pruning. We remove components from randomly weighted neural networks – neurons from fully connected layers and filters from convolutional layers – such that the loss of the networks decreases continuously. This results smaller, dense networks whose accuracy is higher than their initial version.

References

[1] H. Zhou, J. Lan, R. Liu, and J. Yosinski, "Deconstructing lottery tickets: Zeros, signs, and the supermask," in *Advances in Neural Information Processing Systems 32*, pp. 3592–3602, Curran Associates, Inc., 2019.

[2] V. Ramanujan, M. Wortsman, A. Kembhavi, A. Farhadi, and M. Rastegari, "What's hidden in a randomly weighted neural network?," 2019.

[3] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," in *International Conference on Learning Representations*, 2019.