

Reading List

- Han, S., Mao, H. and Dally, W.J., 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. arXiv preprint arXiv:1510.00149.
- Frankle, J. and Carbin, M., 2018. The lottery ticket hypothesis: Finding sparse, trainable neural networks. ICLR 2019
- Malach, E., Yehudai, G., Shalev-Schwartz, S. and Shamir, O., 2020, November. Proving the lottery ticket hypothesis: Pruning is all you need. In International Conference on Machine Learning (pp. 6682-6691). PMLR.
- Frankle, J., Dziugaite, G.K., Roy, D. and Carbin, M., 2020, November. Linear mode connectivity and the lottery ticket hypothesis. In International Conference on Machine Learning (pp. 3259-3269). PMLR.
- Zhou, H., Lan, J., Liu, R. and Yosinski, J., 2019. Deconstructing lottery tickets: Zeros, signs, and the supermask. Advances in neural information processing systems, 32.
- Liu, Z., Sun, M., Zhou, T., Huang, G. and Darrell, T., 2018. Rethinking the value of network pruning. arXiv preprint arXiv:1810.05270.
- Gale, T., Elsen, E. and Hooker, S., 2019. The state of sparsity in deep neural networks. arXiv preprint arXiv:1902.09574.
- Ramanujan, V., Wortsman, M., Kembhavi, A., Farhadi, A. and Rastegari, M., 2020. What's hidden in a randomly weighted neural network?. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 11893-11902).
- Pensia, A., Rajput, S., Nagle, A., Vishwakarma, H. and Papailiopoulos, D., 2020. Optimal lottery tickets via subset sum: Logarithmic over-parameterization is sufficient. Advances in Neural Information Processing Systems, 33, pp.2599-2610.