PhD Research Topic:
Phase Transitions in Artificial Neural Networks

Research axis of the 3IA: 1
Supervisor: Emanuele Natale, CNRS (I3S & INRIA), emanuele.natale@univ-cotedazur.fr
Lab and team: COATI team, I3S & INRIA UCA

Apply by sending an email directly to the supervisor

Project description

It is well known that large ANNs can be trained to achieve very good performance on a variety of tasks, and *once* they have achieved such good performance, they can be *pruned (i.e., sparsified)* to a small fraction of their initial weights without significant loss of performance. It is then natural to ask whether one could *directly train sparse ANNs* that are as sparse as those obtained by first training and then pruning a large ANN. It is empirically observed, however, that trying to train a sparse ANN naively in this way doesn't lead to good performance. The *Lottery Ticket Hypothesis (LTH)* [FC18] establishes that such sparse ANNs (called *lottery tickets*) do exist by empirically showing that, through a relatively computationally expensive procedure, *sufficiently large randomly initialized ANNs contain a sparse subnetwork that can be trained to good performance*.

In 2019 and 2020, a couple of papers [ZLL19,RWK20] showed that it is possible to use gradient descent to *train by pruning*: given a sufficiently large and randomly initialized ANN, one can efficiently learn to *identify a subnetwork* that *performs well* on a given classification task, *without changing to the initial random weights*. Such an empirical observation was very relevant in the context of the LTH: not only do sufficiently large randomly initialized ANNs contain subnetworks that can be efficiently trained, but they also contain subnetworks that already perform well. This motivated subsequent research on a stronger version of the LTH: the *Strong Lottery Ticket Hypothesis* [PRN20,MYS20], which investigates how large the initial random ANN should be in order to be able to approximate *any* function within a given class by appropriate pruning (in other words, one tries to understand how *complete* is the space of possible functions represented by the set of possible subnetworks of the random ANN).

The SLTH has been proven for several ANN architectures [NFG24], with the goal of providing insights on the tradeoff between sparsity and overparameterization, and the limits of ANN compression techniques such as pruning. Current results suffer, however, of two fundamental limitations. First, they exhibit a gap between upper and lower bounds for the SLTH, which hinders practical predictions. Secondly, they mathematically rely on the ANN weights being continuous. Not only are practical ANNs finite precision, but they also heavily rely on quantization techniques.

Research on the SLTH has crucially relied on the fundamental problem of Random Subset Sum (RSS) [L98,DDG23], in which one is asked to prove how large a set of random numbers Ω needs to be so that every number in a given target set Z can be approximated by the sum of a suitable subset of Ω. Interestingly, the RSS is closely related, in a precise way, to the problem of Random Number Partition (RNP), for which sharp analyses in the discrete setting have been provided [BCP01,BCM04]. The purpose of this project is to leverage these classical results on RNP to provide sharp bounds on the SLTH which also account for weight quantization. The latter aspect, in particular, would be a first of its kind and could offer deep insights into a technique which is universally used nowadays for making deep learning viable.

References

BCM04. Borgs, C., Chayes, J. T., Mertens, S. & Pittel, B. Phase diagram for the constrained integer partitioning problem. Random Structures & Algorithms, 2004.
BCP01. Borgs, C., Chayes, J. & Pittel, B. Phase transition and finite-size scaling for the integer partitioning problem. Random Structures & Algorithms, 2001.
DDG23. Da Cunha, A. C. W., d'Amore, F., Giroire, F., Lesfari, H., Natale, E. & Viennot, L. Revisiting the Random Subset Sum Problem. In *ESA 2023.*
FC18. Frankle, J. & Carbin, M. The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks. in *ICLR* 2018.
L98. Lueker, G. S. Exponentially small bounds on the expected optimum of the partition and subset sum problems. *Random Structures & Algorithms* 1998.
MYS20. Malach, E., Yehudai, G., Shalev-shwartz, S. & Shamir, O. Proving the lottery ticket hypothesis: pruning is all you need. In *JMLR 2020*.
NFG24. Natale, E., Ferré, D., Giambartolomei, G., Giroire, F., Mallmann-Trenn, F. On the Sparsity of the Strong Lottery Ticket Hypothesis. In *Neurips 2024*.
PRN20. Pensia, A., Rajput, S., Nagle, A., Vishwakarma, H. & Papailiopoulos, D. Optimal lottery tickets via SUBSETSUM: logarithmic over-parameterization is sufficient. In *Neurips 2020*.
RWK20. Ramanujan, V., Wortsman, M., Kembhavi, A., Farhadi, A. & Rastegari, M. What's Hidden in a Randomly Weighted Neural Network? In *CVPR* 2020.
ZLL19. Zhou, H., Lan, J., Liu, R. & Yosinski, J. Deconstructing Lottery Tickets: Zeros, Signs, and the Supermask. In *NIPS (Neurips)* 2019.