Assegno di Ricerca annuale su fondi FAIR

Title: Mathematical Physics of Learning with Deep Boltzmann Machines

Months: 12

Requirements: Master degree in the areas of Mathematics, Physics or Data Science.

Research Project in brief:

In recent years, Statistical Mechanics (SM) has proven to be fundamental for building a solid theoretical background in machine learning with neural networks, where training is essentially a high-dimensional statistical inference problem. Recently, the self-supervised paradigm has attracted significant attention from the scientific community due to its enormous success in practical applications.

Self-supervised learning (SSL) is a machine learning technique in which a model is trained to use labels generated automatically from the data itself, rather than using manually provided labels (as is the case with supervised learning). This method exploits intrinsic structures or relationships within the data to create learning tasks that enable the model to acquire useful representations. For instance, the model could be trained to predict a missing part of an image or a textual sequence, using the remaining parts as context.

Once the model has learned general representations through self-supervised tasks, these representations can be used for pretraining other specific tasks, such as classification or recognition. This approach falls under the general paradigm of transfer learning.

A relatively simple self-supervised learning architecture is the Deep Boltzmann Machine (DBM), composed of multiple simpler architectures called Restricted Boltzmann Machines stacked on top of one another. In the SM formulation, a DBM is a multi-species spin model where the interactions have to be learned from the data. From the generative model (or direct problem) perspective, DBMs have been extensively studied, but many open problems remain, particularly regarding how to generalize Parisi's theory to non-convex structures like this, from the computation of free energy to the characterization of equilibrium states.

The study of phase diagrams in these models can help answer relevant questions such as:

- How much data, sampled from the environment, is necessary for the machine to efficiently learn a representation of the environment and thus make accurate predictions?
- How does the machine learn the structure of the dataset, and how is this encoded in the interactions?

An insightful approach toward a mathematical formalization of the problem is known as the teacher-student scenario: a neural network with a given architecture (teacher) defines the model for the environment, while another network (student) has to learn something about this environment by leveraging a dataset provided by the teacher. In this controlled setting, it is possible to precisely define both the accuracy in reproducing the training set and the generalization error, and to investigate the properties of the student's representation of the environment as a function of the architecture (unit priors, width, depth, regularization).

The aim of this project is to study the statistical mechanics of RBM and DBM within the teacher-student framework, both in the Nishimori regime, where the student is fully informed about the teacher's machine, and outside of it. We will also explore the case of multiple students trying to learn from the same teacher, either cooperatively or competitively, as this situation allows for the investigation of aspects related to the fine-tuning process and transfer learning in self-supervised learning.

Finanziato dall'Unione Europea - NextGenerationEU a valere sul Piano Nazionale di Ripresa e Resilienza (PNRR) – Missione 4 Istruzione e ricerca – Componente 2 Dalla ricerca all'impresa - Investimento Investimento 1.3 , Avviso D.D. 341 del 15/03/2022, dal titolo: Future Artificial Intelligence Research, codice proposta PE0000013 - CUP J33C22002830006