Mini-courses: Titles & Abstracts

Prof. Elena Agliari *(Sapienza University of Rome)*
Title: Hebbian neural networks

Abstract: This course provides a comprehensive introduction to Hebbian neural networks, exploring their theoretical foundations, paradigmatic models, and applications.

We begin with an overview of attractor neural networks, focusing on the canonical Hopfield model. Restricting to the fairly standard replica-symmetric description, we derive their analytical solutions by Guerra's interpolation methods and summarise their information processing capabilities through suitable phase diagrams.

Next, we examine variations on the theme and compare different synaptic plasticity mechanisms (e.g., pseudo-inverse) beyond the classical Hebbian rule.

Then, we investigate the duality between Hebbian neural networks and Boltzmann Machines and show how the well-established knowledge available for the former can be translated to the latter (e.g., providing recipes for optimal pre-training).

Finally, we explore alternative embedding architectures, moving beyond the fully connected topology, to analyse their impact on storage capacity and the possible emergence of non-trivial computational tasks (e.g., pattern disentanglement).

Monday 13:30-15:30, Tuesday 09:00-12:00, Wednesday 09:00-10:30

Prof. Gérard Ben Arous *(CIMS, New York University)*
Title: The elastic manifold: topological complexity and free energy

Abstract: These talks will be based on joint works with Paul Bourgade (Courant) and Ben Mc Kenna (Georgia Tech) (CPAM 2024), and more recent joint works with Pax Kivimae (Courant). I will present here recent progress on the Elastic Manifold, as introduced by Daniel Fisher in the 80's and studied by Giorgio Parisi and Marc Mézard in the 90's, and by a long line of Physics works including recent ones by Yan Fyodorov and Pierre Le Doussal (2020). This lattice model studies the interaction between the inherent disorder of spin glasses and the taming influence of an elastic interaction.

-We will first cover the question of the topological complexity of this model, following joint works with Paul Bourgade and Ben Mc Kenna, which will bring us to an interesting question of Random Matrix Theory.

-I will then present the more recent work with Pax Kivimae proving a Parisi formula for the quenched free energy, following the "cavity method" as developed by Guerra and Talagrand, and show some consequences for the behavior of the Gibbs measure, in particular for some polymer models.

Thursday 09:00-10:30 and 13:30-15:30, Friday 09:00-10:30 and 11:00-12:00

Prof. Aukosh Jagannath *(University of Waterloo, Canada)*
Title: Introduction to high-dimensional learning dynamics

Abstract: We will give a broad survey of recent progress in understanding
the training dynamics and local geometry of high-dimensonal learning tasks.
We will touch on three distinct but deeply connected topics: sample complexity bounds,
high-dimensional scaling limits, and the local geometry of high-dimensional learning.

Monday 09:30-12:00, Tuesday 13:30-15:30, Thursday 16:00-18:00

Prof. Andrea Montanari *(Stanford University)*
Title: Generalization and overfitting in two-layer neural networks

Abstract: The central question of learning theory is: how is it possible that a model trained on a finite dataset is capable of responding correctly to unseen inputs? I will review the classical answer to this question (based on uniform convergence theory), discuss why this answer is not fully adequate to understand deep learning and overview recent developments.

Tuesday 16:00-18:00, Wednesday 11:00-13:00, Thursday 11:00-12:00, Friday 12:00-13:00