Dipartimento di Matematica e Fisica
"Ennio De Giorgi"

CNR NANOTEC
INSTITUTE OF NANOTECHNOLOGY

Book of Abstract

Interdisciplinary Statistical Physics in Lecce

18th - 19th September 2025

MUR

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

Finanziato dall'Unione europea
NextGenerationEU

UNI SALENTO

Contents

1 Decoding emergent patterns in large microbial communities through the lens of disordered systems

Author: Ada Altieri
Affiliation: Université Paris Cité, Paris

The remarkable biodiversity found in natural ecosystems has recently drawn increasing interest, not only from ecologists but also from theoretical physicists. In this talk, I will present a disordered variant of the Generalized Lotka-Volterra (GLV) model, which incorporates random species interactions and demographic noise, as a proof of concept for its applicability to the human (gut) microbiota. By analyzing metagenomic data from both healthy individuals and patients with inflammatory bowel disease, I will show that the distinct physiological states of the gut microbiome correspond to different noise-driven and disorder-driven regimes in the GLV model [1, 2]. Finally, I will discuss possible extensions of the model to take non-logistic growth behaviors into account, which we propose as stabilizing mechanisms that remain effective even as species richness increases [3].

References:

[1] A. Altieri, F. Roy, C. Cammarota, G. Biroli, Phys. Rev. Lett. 126 (2021)

[2] J. Pasqualini, A. Maritan, A. Rinaldo, S. Facchin, E. V. Savarino, A. Altieri* & S. Suweis*, eLife14:RP105948 (2025)

[3] I. Hatton, O. Mazzarisi, A. Altieri, M. Smerlak, Science 383 (2024)

2 An atlas for matchings

Author: Sergio Caracciolo
Affiliation: Università degli studi di Milano, Milano

We (in collaboration with Gabriele Sicuro and Andrea Sportiello) would like to understand the critical properties of the optimal matching problem. In particular, whether conformality holds. To this aim, we introduce a collection of deformations and excitations in a very general setting and study their structure, uncovering unexpected relations with spanning trees enjoying a series of geometrical properties when two-dimensional Euclidean embeddings are considered.

3 Mutational Paths in Sequence Landscapes inferred from Data

Author: Simona Cocco
Affiliation: École Normale Supérieure, Paris

In the first part of the talk, I will introduce unsupervised generative models learned on protein sequence data, in particular Restricted Boltzmann Machine. I will then introduce a Monte Carlo path sampling algorithm to sample mutational paths in the sequence landscape inferred by our model. I will apply the mutational paths algorithm to sample evolutionary paths for the WW domain, both between two different natural proteins with the same binding specificity and between two natural proteins with different binding specificities. I will then apply our path sampling algorithm and its mean-field derivation to sample evolutionary paths for the Receptor Binding Domain of the Spike protein in SARS-CoV-2 under the pressure of antibody escaping.

4 LLM and Number Theory: a Case Study

Author: Pierluigi Contucci
Affiliation: Università di Bologna, Bologna

We explore a novel representation of the natural numbers through planar rooted trees, where each integer is uniquely mapped to a combinatorial structure via its iterated Euclidean decomposition. This representation allows us to treat the sequence of natural numbers as a text governed by syntactic rules and combinatorial constraints. We investigate this text-like structure using tools from statistical physics and complex systems, such as Zipf's law, correlation analysis, and measures of anomalous diffusion. We test the ability of a Large Language Model (LLM) to capture this structure and to predict factorizations. This is done by implementing a custom tokenizer based on sub-word encoding and training models from scratch under both next-word and masked-prediction tasks. A Hidden Markov Model was also trained on the same dataset to provide a classical benchmark and to motivate the use of deep learning. Our approach offers insight into the learnability and structure of this synthetic text and its formal grammar, revealing unexpected similarities with natural language.

5 RBMs, powerful inference machines

Author: Aurélien Decelle
Affiliation: Universidad Politécnica de Madrid, Madrid

Restricted Boltzmann Machines (RBMs) are unsupervised machine learning models structured as bipartite neural networks. They capture statistical correlations in data through Boltzmann weights derived from an energy function. In this talk, I will demonstrate how analyzing this energy function provides deep and interpretable insights into the data, establishing RBMs as powerful tools for inference. Notably, the RBM's energy function can be mapped onto a model of interacting spins, analogous to those extensively studied in statistical physics. Unlike traditional pairwise interaction models, commonly used in inverse statistical problems (such as maximum entropy approaches), this mapping inherently includes higher-order interactions. Importantly, this framework enables the inference of complex, higher-order interactions without requiring an explosion in the number of adjustable parameters.

I will show how the RBM successfully recovers both pairwise interactions in Ising systems, but also three-body interactions, including when dealing with disordered models. Finally, I will show how the RBM outperforms the Boltzmann Machine in predicting the contacts in the three-dimensional structure using homologous protein family sequences, and in particular that we observe a strong signal that three-body interactions are important in this real dataset.

6 Inferring the Turbulent Breakup of Colloidal Aggregates Using Graph Neural Networks

Author: Alessandra Lanotte
Affiliation: CNR - Nanotec, Lecce

Solid aggregates in dilute turbulent colloidal suspensions can break due to strong local shear fluctuations. In this talk, I will present a data-driven model based on Graph Neural Networks, which have been trained and tested to predict whether an aggregate will break or not, once the turbulent velocity gradient statistics are known.

Acknowledgments. This research is developed under the project "National Centre for HPC, Big Data and Quantum Computing", within the national PNRR plan, financed by NextGenerationEU, and the MUR project PRIN 2022 PNRR P20225AEF4 “BREAKUP - Numerical and Experimental Investigation of Small Plastics Breakup in Complex Flows”.

7 Spurious overlaps in attractor networks

Author: Enzo Marinari
Affiliation: Sapienza Università di Roma, Roma

In attractor neural networks, spurious overlaps are detrimental correlations between the network's state and memories that are distinct from the memory that is being retrieved. We show that while individual spurious overlaps are vanishingly small in the large network size limit, their collective effect can suppress the mean input currents to neurons, promoting sparse neural activity and as a result increasing storage capacity. While classical attractor models do not exhibit this behavior, plasticity rules inferred from neurophysiological experiments do. These findings offer new insights into memory storage in biologically inspired networks.

Work by M. Benedetti, N. Brunel, EM and U. Pereira

8 Unifying principles of learning with and without brains

Author: Matteo Marsili
Affiliation: The Abdus Salam International Centre for Theoretical Physics (ICTP), Trieste

The aim of the talk is to illustrate an approach to learning based on a notion of absolute relevance. I will expand on its relation to inference and a principle-based approach to abstraction, understanding, intelligence, and awareness in simple, yet I hope not totally trivial, terms, with emphasis on thermodynamic efficiency.

9 Hearings, mishearings and word decryption

Author: Anita Mehta
Affiliation: Oxford University, Oxford

We present a model of speech perception, where words correspond to the attractors of a suitably chosen descent dynamics. The resulting lexicon is rich in short words, and much less so in longer ones, as befits a reasonable word-length distribution. We separately examine the decryption of short and long words in the presence of mishearings. In the regime of short words, the algorithm either quickly retrieves a word, or proposes another valid word. In the regime of longer words, the behaviour is markedly different. While the successful decryption of words continues to be relatively fast, there is a finite probability of getting lost permanently, as the algorithm wanders round the landscape of suitable words without ever settling on one.

10 Are Neural Networks Collision Resistant?

Author: Marc Mézard
Affiliation: Bocconi University, Milano

When training neural networks to classify a dataset, one finds a set of weights from which the network produces a label for each data point. We study the algorithmic complexity of finding a collision in a single-layer neural net, where a collision is defined as two distinct sets of weights that assign the same labels to all data. For binary perceptrons with oscillating activation functions, we establish the emergence of an overlap gap property in the space of collisions. This is a topological property believed to be a barrier to the performance of efficient algorithms. The hardness is supported by numerical experiments using approximate message passing algorithms and simulated annealing, for which the algorithms stop working well below the value predicted by our analysis.

Neural networks provide a new category of candidate collision-resistant functions, which for some parameter settings depart from constructions based on lattices. Beyond relevance to cryptography, our work uncovers new forms of computational hardness emerging in large neural networks which may be of independent interest.

11 The lower critical dimension for spin glasses and the behavior of the correlation function in a tube geometry

Author: Giorgio Parisi
Affiliation: Sapienza Università di Roma, Roma

More than thirty years ago, it was predicted that the lower critical dimension of spin glasses is equal to 5/2. This computation implies that the correlation length in a tube geometry behaves as $L^{4/3}$, L being the transverse dimension. This prediction has been successfully compared with large-scale simulations for lattices up to $24^2 \times 72$, with excellent results. The correlation length in the Replica Symmetry Breaking scenario behaves in the same way, proving that Replica symmetry is broken.

We have also simulated a mean-field tube geometry, and we have obtained similar results.

12 Statistical mechanics of extensive-width Bayesian neural networks near interpolation

Author: Mauro Pastore
Affiliation: The Abdus Salam International Centre for Theoretical Physics (ICTP), Trieste

In this contribution, I will present a statistical physics analysis of the supervised learning of a two-layer fully connected network with generic weight distribution and activation function, whose hidden layer is large but remains proportional to the input dimension. We focus on the Bayes-optimal learning in a teacher-student scenario, i.e., with a dataset generated by another network with the same architecture. We operate around interpolation, where the number of trainable parameters and of data are comparable and feature learning emerges. Our analysis uncovers a rich phenomenology with various learning transitions as the number of data increases. In particular, the more strongly the features (i.e., hidden neurons of the target) contribute to the observed responses, the fewer data are needed to learn them. Moreover, when the data is scarce, the model only learns non-linear combinations of the teacher weights, rather than "specialising" by aligning its weights with the teacher's. Specialisation occurs only when enough data becomes available, but it can be hard to find for practical training algorithms, possibly due to statistical-to-computational gaps.

13 Optimizing Spin Glasses with the help of Machine Learning

Author: Federico Ricci-Tersenghi
Affiliation: Sapienza Università di Roma, Roma

Spin glasses are the prototype of models with a complex energy landscape, which is hard to explore. This makes both optimization and sampling problems extremely challenging in this class of models. For this reason, they are often used as benchmarks for new algorithms. Recently, the development of generative models has seen an impressive boost with many diverse applications. It is natural to ask whether these machine learning techniques can help improve optimization and sampling in spin-glass models. I will present the result of a first systematic attempt to find the ground state of spin-glass models, with algorithms that integrate local Monte Carlo moves with global updates generated through an autoregressive model.

14 Molecular Replica Symmetry Breaking

Author: Felix Ritort
Affiliation: Small Biosystems Lab, Condensed Matter Physics Department, University of Barcelona, Carrer Marti i Franques 1, 08028 Barcelona (Spain); Institut de Nanociència i Nanotecnologia (IN2UB), 08028 Barcelona (Spain); Reial Academica de Ciencies i Arts de Barcelona, La Rambla 115, 08002 Barcelona (Spain)

I present an overview of mechanical unzipping experiments of single nucleic acids [1] and proteins [2] showing the power of single-molecule techniques to unravel molecular energy landscapes with kcal/mol accuracy [3,4]. I argue that pulling experiments offer an ideal playground to explore rugged free-energy landscapes and replica symmetry breaking at the single-molecule level [5,6].

References:

[1] P. Rissone and F. Ritort, Nucleic Acid Thermodynamics Derived from Mechanical Unzipping Experiments. Life, 12 (2022) 1089.

[2] M. Rico-Pasto, A. Zaltron, S. Davies, S. Frutos and F. Ritort, Molten-globule like transition state of protein barnase measured with calorimetric force spectroscopy, Proceedings of the National Academy of Sciences 119 (11) e2112382119 (2022)

[3] P. Rissone, M. Rico-Pasto, S. B. Smith and F. Ritort, DNA Calorimetric Force Spectroscopy at Single Base Pair Resolution, Nature Communications 16 (2025) 2688

[4] P. Rissone, C. V. Bizarro, F. Ritort, Stem loop formation drives RNA folding in mechanical unzipping experiments, Proceedings of the National Academy of Sciences 119 (3) e2025575119 (2022)

[5] P. Rissone, A. Severino, I. Pastor and F. Ritort, Universal cold phase transitions in RNA, Proceedings of the National Academy of Sciences 121 (34) e2408313121 (2024)

[6] F. Ritort, Molecular Replica Symmetry Breaking, in Spin-Glass Theory and Far Beyond: Replica Symmetry Breaking after 40 years. Charbonneau, P., Marinari, E., Parisi, G., Ricci-tersenghi, F., Sicuro, G., Zamponi, F., & Mezard, M. (Eds.). Chapter 27, World Scientific, Singapore (2023)

15 Long-term behaviour of learning machines under closed-loop Learning

Author: Yasser Roudi
Affiliation: King's College London, London

In a closed-loop learning set-up, model parameters are repeatedly estimated from data generated from the model itself. Given the possibility that large neural network models may, in the future, be primarily trained with data generated by artificial neural networks themselves, understanding the long-term behaviour of learning machines under closed-loop learning is receiving a lot of attention.

In this talk, we study this process in the case of models that belong to exponential families. We show that in this case, maximum likelihood estimation of the parameters results in a process that converges to absorbing states that amplify initial biases present in the data. This outcome may, however, be prevented if the data contains at least one data point generated from a ground truth model, by relying on maximum a posteriori estimation or by introducing regularisation. Simulations also show that similar results are observed in the Hopfield or Restricted Boltzmann Machines.

16 Probing and Learning the Hidden "Grammar" of Data with Diffusion Models

Author: Antonio Sclocchi
Affiliation: University College London, London

Generative artificial intelligence has revolutionized the modeling of complex, high-dimensional data such as images and text. Its success stems from the ability of deep networks to capture hierarchical and compositional structures—the hidden "grammar" of data, where simple elements combine into larger patterns across scales.

In this talk, I show how diffusion models—generative algorithms that create data by progressively reversing stochastic noise—can both reveal and learn these grammars. Using a theoretical model of hierarchical data, we uncover a critical phase transition in the generative dynamics: as denoising progresses, high-level features emerge abruptly, accompanied by a diverging dynamical correlation length in the observable data. We further show that diffusion models acquire latent grammar rules sequentially in depth, with sample complexity governed by data correlations. We confirm these predictions in text and image models, where we observe universal signatures of hierarchical organization and critical behavior.

These results establish a new connection between generative AI and concepts from statistical physics—such as renormalization and critical phenomena—offering a new lens on how machines learn to create.

17 The quantum AT line in the self-overlap corrected SK model

Author: Simone Warzel
Affiliation: Technical University of Munich, Munich

One striking difference between the phase diagrams of spin-glass models with transverse and longitudinal fields is the disappearance of the glass phase even at zero temperature. In this talk, I will outline a rigorous argument for the shape of the quantum Almeida-Thouless (AT) line in the self-overlap-corrected version of the Sherrington-Kirkpatrick (SK) model. I will also argue why this helps shed some light on the quantum AT line in the original SK model.

Based on joint work with C. Manai.

18 Neural Network architectures for efficient sampling of statistical physics models

Author: Francesco Zamponi
Affiliation: Sapienza Università di Roma, Roma

The task of sampling efficiently the Gibbs-Boltzmann distribution of disordered systems is important both for the theoretical understanding of these models and for the solution of practical optimization problems. Unfortunately, this task is known to be hard, especially for spin glasses at low temperatures. Recently, many attempts have been made to tackle the problem by mixing classical Monte Carlo schemes with newly devised neural networks that learn to propose smart moves. In this talk I will review a few physically-interpretable deep architectures, and in particular one whose number of parameters scales linearly with the size of the system and that can be applied to a large variety of topologies. I will show that these architectures can accurately learn the Gibbs-Boltzmann distribution for the two-dimensional and three-dimensional Edwards-Anderson models, and specifically for some of its most difficult instances. I will show that the performance increases with the number of layers, in a way that clearly connects to the correlation length of the system, thus providing a simple and interpretable criterion to choose the optimal depth. Finally, I will discuss the performances of these architectures in proposing smart Monte Carlo moves and compare them to state-of-the-art algorithms.

References:

Nearest - Neighbours Neural Network architecture for efficient sampling of statistical physics models, L.M.Del Bono, F.Ricci-Tersenghi, F.Zamponi - Machine Learning: Science and Technology 6, 025029 (2025)

19 Dynamical Learning in Deep Asymmetric Recurrent Neural Networks

Author: Riccardo Zecchina
Affiliation: Bocconi University, Milano

We show that asymmetric deep recurrent neural networks, enhanced with additional sparse excitatory couplings, give rise to an exponentially large, connected cluster of stable internal representations which can be found by different algorithms, including simple iterative dynamics. Building on the geometrical properties of the stable configurations, we propose a distributed learning scheme in which input-output associations emerge naturally from the recurrent dynamics, without any need for gradient evaluation. A critical feature enabling the learning process is the stability of the configurations reached at convergence, even after the removal of the supervisory output signal. Extensive simulations demonstrate that this approach performs competitively on standard AI benchmarks. The model can be generalized in multiple directions, both computational and biological, potentially contributing to narrowing the gap between AI and computational neuroscience.

Funded by:
PRIN 2022 grant Statistical Mechanics of Learning Machines number 20229T9EAT funded by the Italian Ministry of University and Research (MUR) in the framework of European Union - Next Generation EU.
PRIN 2022 PNRR grant Stochastic Modeling of Compound Events (SLIDE) number P2022KZJTZ funded by the Italian Ministry of University and Research (MUR) in the framework of European Union - Next Generation EU.