Probing the energy landscape of neural networks

Carlo Lucibello*[1]

[1]Politecnico di Torino [Torino] (Polito) – Politecnico di Torino - Corso Duca degli Abruzzi, 24 10129 Torino, Italy

Abstract

Training neural networks with very low precision synapses has long been considered a challenging task even for the simplest neural architectures. In this talk I'll present a series of results which emerged from a large-deviation analysis using tools from Statistical Physics, which show that the training problem can be made algorithmically very simple by maximizing a "local entropy": explicitly seeking extensive regions in the space of configurations with low energy. Such regions also have some highly desirable properties, in particular very good generalization capabilities. These results appear to be rather general with respect to the details of the underlying model and of the data, and may be relevant biologically and technologically, as well as apply to other inference and constraint satisfaction problems.

Bibliography

C. Baldassi, C. Borgs, J. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina.

"Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes", PNAS (2016)

*Speaker